Distributed Computing manuscript No.
(will be inserted by the editor)

Toward More Localized Local Algorithms: Removing Assumptions Concerning Global Knowledge

Amos Korman · Jean-Sébastien Sereni · Laurent Viennot

arXiv:1512.03306v1 [cs.DC] 10 Dec 2015

Received: date / Accepted: date

Abstract Numerous sophisticated local algorithm were suggested in the literature for various fundamental problems. Notable examples are the MIS and $(\Delta+1)$-coloring algorithms by Barenboim and Elkin [6], by Kuhn [22], and by Panconesi and Srinivasan [34], as well as the $O(\Delta^2)$-coloring algorithm by Linial [28]. Unfortunately, most known local algorithms (including, in particular, the aforementioned algorithms) are *non-uniform*, that is, local algorithms generally use good estimations of one or more global parameters of the network, e.g., the maximum degree Δ or the number of nodes n.

This paper provides a method for transforming a non-uniform local algorithm into a *uniform* one. Furthermore, the resulting algorithm enjoys the same asymptotic running time as the original non-uniform algorithm. Our method applies to a wide family of both deterministic and randomized algorithms. Specifically, it applies to almost all state of the art non-uniform algorithms for MIS and Maximal Matching, as well as to many results concerning the coloring problem. (In particular, it applies to all aforementioned algorithms.)

To obtain our transformations we introduce a new distributed tool called *pruning algorithms*, which we believe may be of independent interest.

Keywords distributed algorithm · global knowledge · parameters · MIS · coloring · maximal matching

Amos Korman is supported in part by a France-Israel cooperation grant ("Mutli-Computing" project) from the France Ministry of Science and Israel Ministry of Science, by the ANR projects ALADDIN and PROSE, and by the INRIA project GANG. Jean-Sébastien Sereni is partially supported by the French *Agence Nationale de la Recherche* under reference ANR 10 JCJC 0204 01. Laurent Viennot is supported by the european STREP project EULER, and the INRIA project-team GANG.

Amos Korman
CNRS and University Paris Diderot
LIAFA Case 7014
Université Paris Diderot – Paris 7
F-75205 Paris Cedex 13, France.
Tel.: +33-1-57-27-92-56
Fax: +33-1-57-27-94-09
E-mail: Amos.Korman@liafa.jussieu.fr

Jean-Sébastien Sereni
CNRS (LIAFA, Université Denis Diderot), Paris, France
and Department of Applied Mathematics (KAM), Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic
E-mail: sereni@kam.mff.cuni.cz

Laurent Viennot
INRIA and University Paris Diderot
LIAFA Case 7014
Université Paris Diderot – Paris 7
F-75205 Paris Cedex 13, France.
E-mail: Laurent.Viennot@inria.fr

1 Introduction

1.1 Background and Motivation

Distributed computing concerns environments in which many processors, located at different sites, must collaborate in order to achieve some global task. One of the main themes in distributed network algorithms concerns the question of how to cope with *locality* constrains, that is, the lack of knowledge about the global structure of the network (cf., [35]). On the one hand, information about the global structure may not always be accessible to individual processors and the cost of computing it from scratch may overshadow the cost of the algorithm using it. On the other hand, global knowledge is not always essential, and many seemingly global tasks can be efficiently achieved by letting processors

know more about their immediate neighborhoods and less about the rest of the network.

A standard model for capturing the essence of locality is the $\mathcal{LOCAL}$ model (cf., [35]). In this model, the network is modeled by a graph G, where the nodes of G represent the processors and the edges represent the communication links. To perform a task, nodes are woken up simultaneously, and computation proceeds in fault-free synchronous rounds during which every node exchanges messages with its neighbors, and performs arbitrary computations on its data. Since many tasks cannot be solved distributively in an anonymous network in a deterministic way, symmetry breaking must be addressed. Arguably, there are two typical ways to address this issue: the first one is to use randomized algorithms, while the second one is to assume that each node v in the network is initially provided a unique identity $\mathrm{Id}(v)$. A *local algorithm* operating in such a setting must return an output at each node such that the collection of outputs satisfies the required task. For example, a *Maximal Independent Set* (MIS) of a graph G is a set S of nodes of G such that every node not in S has a neighbor in S and no two nodes of S are adjacent. In a local algorithm for the MIS problem, the output at each node v is a bit $b(v)$ indicating whether v belongs to a selected set S of nodes, and it is required that S forms a MIS of G. The *running time* of a local algorithm is the number of rounds needed for the algorithm to complete its operation at each node, taken in the worst case scenario. This is typically evaluated with respect to some parameters of the underlying graph. The common parameters used are the number of nodes n in the graph and the maximum degree Δ of a node in the graph.

To ease the computation, it is often assumed that some kind of knowledge about the global network is provided to each node *a priori*. A typical example of such knowledge is the number of nodes n in the network. It turns out that in some cases, this (common) assumption can give a lot of power to the distributed algorithm. This was observed by Fraigniaud *et al.* [16] in the context of local decision: they introduced the complexity class of decision problems NLD, which contains all decision problems that can be verified in constant time with the aid of a certificate. They proved that, although there exist decision problems that do not belong to NLD, every (computable) decision problem falls in NLD if it is assumed that each node is given the value of n as an input.

In general, the amount and type of such information may have a profound effect on the design of the distributed algorithm. Obviously, if the whole graph is contained in the input of each node, then the distributed algorithm can be reduced to a central one. In fact, the whole area of *computation with advice* [9,12, 13,14,15,20,21] is dedicated to studying the amount of information contained in the inputs of the nodes and its effect on the performances of the distributed algorithm. For instance, Fraigniaud *et al.* [15] showed that if each node is provided with only a constant number of bits then one can locally construct a BFS-tree in constant time, and can locally construct a MST in $O(\log n)$ time, while both tasks require diameter time if no knowledge is assumed. As another example, Cohen *et al.* [9] proved that $O(1)$ bits, judiciously chosen at each node, can allow a finite automaton to distributively explore every graph. As a matter of fact, from a radical point of view, for many questions (e.g., MIS and Maximal Matching), additional information may push the question at hand into absurdity: even a constant number of bits of additional information per node is enough to compute a solution—simply let the additional information encode the solution!

When dealing with locality issues, it is desired that the amount of information regarding the whole network contained in the inputs of the nodes is minimized. A local algorithm that assumes that each node is initially given merely its own identity is often called *uniform*. Unfortunately, there are only few local algorithms in the literature that are uniform (e.g., [11,26,29,30,37]). In contrast, most known local algorithms assume that the inputs of all nodes contain upper bounds on the values of some global parameters of the network. Moreover, it is often assumed that all inputs contain the same upper bounds on the global parameters. Furthermore, typically, not only the correct operation of the algorithm requires that upper bounds be contained in the inputs of all nodes, but also the running time of the algorithm is actually a function of the upper bound estimations and not of the actual values of the parameters. Hence, it is desired that the upper bounds contained in the inputs are not significantly larger than the real values of the parameters.

Some attempts to transform a non-uniform local algorithm into a uniform one were made by examining the details of the algorithm at hand and modifying it appropriately. For example, Barenboim and Elkin [6] first gave a non-uniform MIS algorithm for the family of graphs with arboricity $a = O(\log^{1/2-\delta} n)$, for any constant $\delta \in (0, 1/2)$, running in time $O(\log n/ \log\log n)$. (The *arboricity* of a graph being the smallest number of acyclic subgraphs that together contain all the edges of the graph.) At the cost of increasing the running time to $O(\frac{\log n}{\log\log n} \log^* n)$, the authors show how to modify their algorithm so that the value of a need not be part of the inputs of nodes. In addition to the MIS algo-

rithms, the work of [6] also contains algorithms that do not require the knowledge of the arboricity, but have the same asymptotic running time as the ones that require it. For example, this corresponds to algorithms computing forests-decomposition and $O(a)$-coloring. Nevertheless, all their algorithms still require the inputs of all nodes to contain a common upper bound on n.

We present general methods for transforming a non-uniform local algorithm into a uniform one without increasing the asymptotic running time of the original algorithm. Our method applies to a wide family of both deterministic and randomized algorithms. In particular, our method applies to all state of the art non-uniform algorithms for MIS and Maximal Matching, as well as to several of the best known results for $(\Delta+1)$-coloring.

Our transformations are obtained using a new type of local algorithms termed *pruning algorithms*. Informally, the basic property of a pruning algorithm is that it allows one to iteratively apply a sequence of local algorithms (whose output may not form a correct global solution) one after the other, in a way that "always progresses" toward a solution. In a sense, a pruning algorithm is a combination of a gluing mechanism and a *local checking* algorithm (cf., [16,32]). A local checking algorithm for a problem Π runs on graphs with an output value at each node (and possibly an input too), and can locally detect whether the output is "legal" with respect to Π. That is, if the instance is not legal then at least one node detects this, and raises an alarm. (For example, a local checking algorithm for MIS is trivial: each node in the set S, which is suspected to be a MIS, checks that none of its neighbors belongs to S, and each node not in S checks that at least one of its neighbors belongs to S. If the check fails, then the node raises an alarm.) A pruning algorithm needs to satisfy an additional *gluing* property not required by local checking algorithms. Specifically, if the instance is not legal, then the pruning algorithm must carefully choose the nodes raising the alarm (and possibly modify their input too), so that a solution for the subgraph induced by those alarming nodes can be well glued to the previous output of the non-alarming nodes, in a way such that the combined output is a solution to the problem for the whole initial graph.

We believe that this new type of algorithms may be of independent interest. Indeed, as we show, pruning algorithms have several types of other applications in the theory of local computation, besides the aforementioned issue of designing uniform algorithms. Specifically, they can be used also to transform a local Monte-Carlo algorithm into a Las Vegas one, as well as to obtain an algorithm that runs in the minimum running time of a given (finite) set of uniform algorithms.

1.2 Previous Work

MIS and coloring: There is a long line of research concerning the two related problems of $(\Delta+1)$-coloring and MIS [3,10,17,18,23,24,28]. A *k-coloring* of a graph is an assignment of an integer in $\{1,\ldots,k\}$ to each node such that no two adjacent vertices are assigned the same integer. Recently, Barenboim and Elkin [4] and independently Kuhn [22] presented two elegant $(\Delta+1)$-coloring and MIS algorithms running in $O(\Delta+\log^* n)$ time on general graphs. This is the best currently-known bound for these problems on low degree graphs. For graphs with a large maximum degree Δ, the best bound is due to Panconesi and Srinivasan [34], who devised an algorithm running in $2^{O(\sqrt{\log n})}$ time. The aforementioned algorithms are not uniform. Specifically, all three algorithms require that the inputs of all nodes contain a common upper bound on n and the first two also require a common upper bound on Δ.

For bounded-independence graphs, Schneider and Wattenhofer [37] designed uniform deterministic MIS and $(\Delta+1)$-coloring algorithms running in $O(\log^* n)$ time. Barenboim and Elkin [6] devised a deterministic algorithm for the MIS problem on graphs of bounded arboricity that requires time $O(\log n/\log\log n)$. More specifically, for graphs with arboricity $a = o(\sqrt{\log n})$, they show that a MIS can be computed deterministically in $o(\log n)$ time, and whenever $a = O(\log^{1/2-\delta} n)$ for some constant $\delta \in (0\,,1/2)$, the same algorithm runs in time $O(\log n/\log\log n)$. At the cost of increasing the running time by a multiplicative factor of $O(\log^* n)$, the authors show how to modify their algorithm so that the value of a need not be part of the inputs of nodes. Nevertheless, all their algorithms require the inputs of all nodes to contain a common upper bound on the value of n. Another MIS algorithm which is efficient for graphs with low arboricity was devised by Barenboim and Elkin [5]; this algorithm runs in time $O(a+a^\epsilon \log n)$ for arbitrary constant $\epsilon > 0$.

Concerning the problem of coloring with more than $\Delta+1$ colors, Linial [27,28], and subsequently Szegedy and Vishwanathan [38], described $O(\Delta^2)$-coloring algorithms with running time $\theta(\log^* n)$. Barenboim and Elkin [4] and, independently, Kuhn [22] generalized this by presenting a tradeoff between the running time and the number of colors: they devised a $\lambda(\Delta+1)$-coloring algorithm with running time $O(\Delta/\lambda+\log^* n)$, for any $\lambda \geqslant 1$. All these algorithms require the inputs of all nodes to contain common upper bounds on both n and Δ.

Barenboim and Elkin [5] devised a $\Delta^{1+o(1)}$ coloring algorithm running in time $O(f(\Delta)\log\Delta\log n)$, for an arbitrarily slow-growing function $f = \omega(1)$. They

Problem	Parameters	Time	Ref.	This paper (uniform)	Corollary 1
Det. MIS and (Δ+1)-coloring	n, Δ	$O(\Delta + \log^* n)$	[4,22]	$\min\left\{O(\Delta + \log^* n), 2^{O(\sqrt{\log n})}\right\}$	(i)
	n	$2^{O(\sqrt{\log n})}$	[34]		(ii)
Det. MIS (arboricity $a = o(\sqrt{\log n})$)	n, a	$o(\log n)$	[6]	$o(\log n)$	(i)
Det. MIS (arboricity $a = O(\log^{1/2-\delta} n)$)	n, a	$O(\log n/\log\log n)$	[6]	$O(\log n/\log\log n)$	(i)
Det. $\lambda(\Delta+1)$-coloring	n, Δ	$O(\Delta/\lambda + \log^* n)$	[4,22]	$O(\Delta/\lambda + \log^* n)$	(iii)
Det. $O(\Delta)$-edge-coloring	n, Δ	$O(\Delta^\epsilon + \log^* n)$	[7]	$O(\Delta^\epsilon + \log^* n)$	(v)
Det. $O(\Delta^{1+\epsilon})$-edge-coloring	n, Δ	$O(\log\Delta + \log^* n)$	[7]	$O(\log\Delta + \log^* n)$	(v)
Det. Maximal Matching	n or Δ	$O(\log^4 n)$	[19]	$O(\log^4 n)$	(vi)
Rand. $(2, 2(c+1))$-ruling set	n	$O(2^c \log^{1/c} n)$	[36]	$O(2^c \log^{1/c} n)$	(vii)
Rand. MIS	uniform	$O(\log n)$	[1,30]		

Table 1 Comparison of $\mathcal{LOCAL}$ algorithms with respect to the use of global parameters. "Det." stands for deterministic, and "Rand." for randomized.

also produced an $O(\Delta^{1+\epsilon})$-coloring algorithm running in $O(\log\Delta\log n)$-time, for an arbitrarily small constant $\epsilon > 0$, and an $O(\Delta)$-coloring algorithm running in $O(\Delta^\epsilon \log n)$ time, for an arbitrarily small constant $\epsilon > 0$. All these coloring algorithms require the inputs of all nodes to contain the values of both Δ and n. Other deterministic non-uniform coloring algorithms with number of colors and running time corresponding to the arboricity parameter were given by Barenboim and Elkin [5, 6].

Efficient deterministic algorithms for the edge-coloring problem can be found in several papers [5,7,33]. In particular, Panconesi and Rizzi [33] designed a simple deterministic local algorithm that finds a $(2\Delta-1)$-edge-coloring of a graph in time $O(\Delta + \log^* n)$. Recently, Barenboim and Elkin [7], designed an $O(\Delta)$-edge-coloring algorithm running in time $O(\Delta^\epsilon)+\log^* n$, for any $\epsilon > 0$, and an $O(\Delta^{1+\epsilon})$-edge-coloring algorithm running in time $O(\log\Delta) + \log^* n$, for any $\epsilon > 0$. All these algorithms require the inputs of all nodes to contain common upper bounds on both n and Δ.

Randomized algorithms for MIS and (Δ+1)-coloring running in expected time $O(\log n)$ were initially given by Luby [30] and, independently, by Alon *et al.* [1].

Recently, Schneider and Wattenhofer [36] constructed the best currently-known non-uniform $(\Delta+1)$-coloring algorithm, which runs in time $O(\log\Delta + \sqrt{\log n})$. They also provided random algorithms for coloring using more colors. For every positive integer c, a randomized algorithm for $(2, 2(c+1))$-ruling set running in time $O(2^c \log^{1/c} n)$ is also presented. (A set S of nodes in a graph being (α,β)-*ruling* if every node not in S is at distance at most β of a node in S and no two nodes in S are at distance less than α.) All these algorithms of Schneider and Wattenhoffer [36] are not uniform and require the inputs of all nodes to contain a common upper bound on n.

Maximal Matching: A *maximal matching* of a graph G is a set M of edges of G such that every edge not in M is incident to an edge in M and no two edges in M are incident. Schneider and Wattenhofer [37] designed a uniform deterministic maximal matching algorithm on bounded-independence graphs running in $O(\log^* n)$ time. For general graphs, however, the state of the art maximal matching algorithm is not uniform: Hanckowiak *et al.* [19] presented a non-uniform deterministic algorithm for maximal matching running in time $O(\log^4 n)$. This algorithm assumes that the inputs of all nodes contain a common upper bound on n (this assumption can be omitted for some parts of the algorithm under the condition that the inputs of all nodes contain the value of Δ).

1.3 Our Results

The main conceptual contribution of the paper is the introduction of a new type of algorithms called *pruning algorithms*. Informally, the fundamental property of this type of algorithms is to allow one to iteratively run a sequence of algorithms (whose output may not necessarily be correct everywhere) so that the global output does not deteriorate, and it always progresses toward a solution.

Our main application for pruning algorithm concerns the problem of locally computing a global solution while minimizing the necessary global information contained in the inputs of the nodes. Addressing this, we provide a method for transforming a non-uniform local algorithm into a uniform one without increasing the asymptotic running time of the original algorithm. Our method applies to a wide family of both deterministic and randomized algorithms; in particular, it applies to many of the best known results concerning classical problems such as MIS, Coloring, and Maximal Matching. (See Table 1.2 for a summary of some of the uni-

form algorithms we obtain and the corresponding state of the art existing non-uniform algorithms.)

In another application, we show how to transform a Monte-Carlo local algorithm into a Las Vegas one. Finally, given a constant number of uniform algorithms for the same problem whose running times depend on different parameters—which are unknown to nodes—we show a method for constructing a uniform algorithm solving the problem, that on every instance runs asymptotically as fast as the fastest algorithm among those given algorithms.

Stating our main results requires a number of formal definitions, so we defer the precise statements to later parts of the paper. Rather, we provide here some interesting corollaries of our results. References for the corresponding non-uniform algorithms are provided in Table 1.2. (The notion of "moderately-slow function" used in item (iii) below is defined in Section 2.)

Corollary 1

(i) *There exists a uniform deterministic algorithm solving* `MIS` *on general graphs in time*

$$\min\{g(n), h(\Delta, n), f(a, n)\},$$

where $g(n) = 2^{O(\sqrt{\log n})}$, $h(\Delta, n) = O(\Delta + \log^* n)$, *and* $f(a, n)$ *is bounded as follows.* $f(a, n) = o(\log n)$ *for graphs of arboricity* $a = o(\sqrt{\log n})$, $f(a, n) = O(\log n / \log\log n)$ *for arboricity* $a = O(\log^{1/2-\delta} n)$, *for some constant* $\delta \in (0\,, 1/2)$*; and otherwise:* $f(a, n) = O(a + a^\epsilon \log n)$, *for arbitrary small constant* $\epsilon > 0$.

(ii) *There exists a uniform deterministic algorithm solving the* $(\Delta+1)$*-coloring problem on general graphs in time* $\min\{O(\Delta + \log^* n), 2^{O(\sqrt{\log n})}\}$.

(iii) *There exists a uniform deterministic algorithm solving the* $\lambda(\Delta+1)$*-coloring problem on general graphs and running in time* $O(\Delta/\lambda + \log^* n)$, *for any* $\lambda \geqslant 1$, *such that* Δ/λ *is a moderately-slow function. In particular, there exists a uniform deterministic algorithm solving the* $O(\Delta^2)$*-coloring problem in time* $O(\log^* n)$.

(iv) *The following uniform deterministic coloring algorithms exist.*

- *A uniform* $\Delta^{1+o(1)}$*-coloring algorithm running in time* $O(f(\Delta)\log\Delta\log n\log\log n)$, *for an arbitrarily slow-growing function* $f = \omega(1)$.
- *A uniform* $O(\Delta^{1+\epsilon})$*-coloring algorithm running in* $O(\log\Delta\log n\log\log n)$ *time, for any constant* $\epsilon > 0$.
- *A uniform* $O(\Delta)$*-coloring algorithm running in* $O(\Delta^\epsilon \log n\log\log n)$ *time, for any constant* $\epsilon > 0$.

(v) - *There exists a uniform deterministic* $O(\Delta)$*-edge-coloring algorithm for general graphs running in time* $O(\Delta^\epsilon + \log^* n)$, *for any constant* $\epsilon > 0$.
- *There exists a uniform deterministic* $O(\Delta^{1+\epsilon})$*-edge-coloring algorithm for general graphs that runs in time* $O(\log\Delta + \log^* n)$, *for any constant* $\epsilon > 0$.

(vi) *There exists a uniform deterministic algorithm solving the maximal matching problem in time* $O(\log^4 n)$.

(vii) *For a constant integer* $c \geqslant 1$, *there exists a uniform randomized algorithm solving the* $(2, 2(c+1))$*-ruling set problem in time* $O(2^c \log^{1/c} n)$.

2 Preliminaries

General definitions: For two integers a and b, we let $[a\,, b] = \{a, a+1, \ldots, b\}$. A vector $\underline{x} \in \mathbf{R}^\ell$ is said to *dominate* a vector $\underline{y} \in \mathbf{R}^\ell$ if $\underline{x}$ is coordinate-wise greater than or equal to $\underline{y}$, that is, $\underline{x}_k \geqslant \underline{y}_k$ for each $k \in [1\,, \ell]$.

For a graph G, we let $V(G)$ and $E(G)$ be the sets of nodes and edges of G, respectively. (Unless mentioned otherwise, we consider only undirected and unweighted graphs.) The *degree* $\deg_G(v)$ of a node $v \in V(G)$ is the number of neighbors of v in G. The *maximum degree of* G is $\Delta_G = \max\{\deg_G(v) : v \in V(G)\}$.

Let u and v be two nodes of G. The *distance* $\mathrm{dist}_G(u, v)$ between u and v is the number of edges on a shortest path connecting them. Given an integer $r \geqslant 0$, the *ball* of radius r around u is the subgraph $B_G(u, r)$ of G induced by the collection of nodes at distance at most r from u. The *neighborhood* $N_G(u)$ of u is the set of neighbors of u, i.e., $N_G(u) = B_G(u, 1) \setminus \{u\}$. In what follows, we may omit the subscript G from the previous notations when there is no risk of confusion.

Functions: A function $f\colon \mathbf{R}^\ell \to \mathbf{R}$ is *non-decreasing* if for every two vectors $\underline{x}$ and $\underline{y}$ such that $\underline{x}$ dominates $\underline{y}$,

$$f(\underline{y}) \leqslant f(\underline{x}).$$

A function $f\colon \mathbf{R}^+ \to \mathbf{R}^+$ is *moderately-slow* if it is non-decreasing and there exists a positive integer α such that

$$\forall i \in \mathbf{N} \setminus \{1\}, \quad \alpha \cdot f(i) \geqslant f(2i).$$

In other words, $f(c \cdot i) = O(f(i))$ for every constant c and every integer i, where the constant hidden in the O notation depends only on c. An example of a moderately-slow function is given by the logarithm.

A function $f\colon \mathbf{R}^+ \to \mathbf{R}^+$ is *moderately-increasing* if it is non-decreasing and there exists a positive integer α such that

$$\forall i \in \mathbf{N} \setminus \{1\}, \quad f(\alpha \cdot i) \geqslant 2f(i) \quad \text{and} \quad \alpha \cdot f(i) \geqslant f(2i).$$

Note that $f(x) = x^{k_1} \log^{k_2}(x)$ is a moderately-increasing function for every two reals $k_1 \geqslant 1$ and $k_2 \geqslant 0$. Moreover, every moderately-increasing function is moderately-slow. On the other hand, some functions (such as the constant functions or the logarithm) are moderately-slow but not moderately-increasing.

A function $f\colon \mathbf{R}^+ \to \mathbf{R}^+$ is *moderately-fast* if it is moderately-increasing and there exists a polynomial P such that

$$\forall x \in \mathbf{R}^+, \quad x < f(x) < P(x).$$

A function $f\colon \mathbf{R}^+ \to \mathbf{R}^+$ *tends to infinity* if

$$\limsup_{x \to \infty} f(x) = \infty,$$

and f is *ascending* if it is non-decreasing and it tends to infinity. (Note that in this case $\lim_{x\to\infty} f(x) = \infty$.)

A function $f\colon (\mathbf{R}^+)^\ell \to \mathbf{R}^+$ is *additive* if there exist ℓ ascending functions $f_1, \ldots, f_\ell$ such that

$$f(x_1, \ldots, x_\ell) = \sum_{i=1}^{\ell} f_i(x_i).$$

Problems and instances: Given a set V of nodes, a *vector for* V is an assignment $\mathbf{x}$ of a bit string $\mathbf{x}(v)$ to each $v \in V$, i.e., $\mathbf{x}$ is a function $\mathbf{x} : V \to \{0,1\}^*$. A *problem* is defined by a collection of triplets: $\Pi = \{(G, \mathbf{x}, \mathbf{y})\}$, where G is a (not necessarily connected) graph, and $\mathbf{x}$ and $\mathbf{y}$ are *input* and *output* vectors for V, respectively. We consider only problems that are closed under disjoint union, i.e., if G_1 and G_2 are two vertex disjoint graphs and $(G_1, \mathbf{x}_1, \mathbf{y}_1), (G_2, \mathbf{x}_2, \mathbf{y}_2) \in \Pi$ then $(G, \mathbf{x}, \mathbf{y}) \in \Pi$, where $G = G_1 \cup G_2$, $\mathbf{x} = \mathbf{x}_1 \cup \mathbf{x}_2$ and $\mathbf{y} = \mathbf{y}_1 \cup \mathbf{y}_2$.

An *instance*, with respect to a given problem Π, is a pair $(G, \mathbf{x})$ for which there exists an output vector $\mathbf{y}$ satisfying $(G, \mathbf{x}, \mathbf{y}) \in \Pi$. In what follows, whenever we consider a collection $\mathcal{F}$ of instances, we always assume that $\mathcal{F}$ is closed under inclusion. That is, if $(G, \mathbf{x}) \in \mathcal{F}$ and $(G', \mathbf{x}') \subseteq (G, \mathbf{x})$ (i.e., G' is a subgraph of G and $\mathbf{x}'$ is the input vector $\mathbf{x}$ restricted to $V(G')$) then $(G', \mathbf{x}') \in \mathcal{F}$. Informally, given a problem Π and a collection of instances $\mathcal{F}$, the goal is to design an efficient distributed algorithm that takes an instance $(G, \mathbf{x}) \in \mathcal{F}$ as input, and produces an output vector $\mathbf{y}$ satisfying $(G, \mathbf{x}, \mathbf{y}) \in \Pi$. The reason to require problems to be closed under disjoint union is that a distributed algorithm operating on an instance $(G, \mathbf{x})$ runs separately and independently on each connected component of G. Let $\mathcal{G}$ be a family of graphs closed under inclusion. We define $\mathcal{F}(\mathcal{G})$ to be $\{G\} \times \{0,1\}^*$.

We assume that each node $v \in V$ is provided with a unique integer referred to as the *identity* of v, and denoted $\mathrm{Id}(v)$; by unique identities, we mean that $\mathrm{Id}(u) \neq \mathrm{Id}(v)$ for every two distinct nodes u and v. For ease of exposition, we consider the identity of a node to be part of its input.

We consider classical problems such as coloring, maximal matching (`MM`), Maximal Independent Set (`MIS`) and the (α, β)-ruling set problem. Informally, viewing the output of a node as a *color*, the requirement of *coloring* is that the colors of two neighboring nodes must be different. In the (α, β)-ruling set problem, the output at each node is Boolean, and indicates whether the node belongs to a set S that must form an (α, β)-ruling set. That is, the set S of selected nodes must satisfy: (1) two nodes that belong to S must be at distance at least α from each other, and (2) if a node does not belong to S, then there is a node in the set at distance at most β from it. `MIS` is a special case of the ruling set problem, specifically, `MIS` is precisely $(2,1)$-ruling set. Finally, given a triplet $(G, \mathbf{x}, \mathbf{y})$, two nodes u and v are said to be *matched* if $(u, v) \in E$, $\mathbf{y}(u) = \mathbf{y}(v)$ and $\mathbf{y}(w) \neq \mathbf{y}(u)$ for every $w \in (N_G(u) \cup N_G(v)) \setminus \{u, v\}$. Thus, the `MM` problem requires that each node u is either matched to one of its neighbors or that every neighbor v of u is matched to one of v's neighbors.

Parameters: Fix a problem Π and let $\mathcal{F}$ be a collection of instances for Π. A *parameter* $\mathtt{p}$ is a positive valued function $\mathtt{p} : \mathcal{F} \to \mathbf{N}$. The parameter $\mathtt{p}$ is *non-decreasing*, if $\mathtt{p}(G', \mathbf{x}') \leqslant \mathtt{p}(G, \mathbf{x})$ whenever $(G', \mathbf{x}') \in \mathcal{F}$ and $(G', \mathbf{x}') \subseteq (G, \mathbf{x})$.

Let $\mathcal{F}$ be a collection of instances. A parameter $\mathtt{p}$ for $\mathcal{F}$ is a *graph-parameter* if $\mathtt{p}$ is independent of the input, that is, if $\mathtt{p}(G, \mathbf{x}) = \mathtt{p}(G, \mathbf{x}')$ for every two instances $(G, \mathbf{x}), (G, \mathbf{x}') \in \mathcal{F}$ such that the input assignments $\mathbf{x}$ and $\mathbf{x}'$ preserve the identities, i.e., the inputs $\mathbf{x}(v)$ and $\mathbf{x}'(v)$ contain the same identity $\mathrm{Id}(v)$ for every $v \in V(G)$. In what follows, we will consider only non-decreasing graph-parameters (note, not all graph-parameters are non-decreasing, an example being the diameter of a graph). More precisely, we will primarily focus on the following non-decreasing graph-parameters: the number n of nodes of the graph G, i.e., $|V(G)|$, the maximum degree $\Delta = \Delta(G)$ of G, i.e., $\max\{\deg_G(u) : u \in V(G)\}$,

and the arboricity $a = a(G)$ of G, i.e., the least number of acyclic subgraphs of G whose union is G.

Local algorithms: Consider a problem Π and a collection of instances $\mathcal{F}$ for Π. An algorithm for Π and $\mathcal{F}$ takes as input an instance $(G, \mathbf{x}) \in \mathcal{F}$ and must terminate with an output vector $\mathbf{y}$ such that $(G, \mathbf{x}, \mathbf{y}) \in \Pi$. We consider the $\mathcal{LOCAL}$ model (cf., [35]). During the execution of a *local* algorithm $\mathcal{A}$, all processors are woken up simultaneously and computation proceeds in fault-free synchronous rounds. In each round, every node may send messages of unrestricted size to its neighbors and may perform arbitrary computations on its data. A message that is sent in a round r, arrives at its destination before the next round $r + 1$ starts. It must be guaranteed that after a finite number of rounds, each node v terminates by writing some final output value $\mathbf{y}(v)$ in its designated output variable (informally, this means that we may assume that a node "knows" that its output is indeed its final output.) The algorithm $\mathcal{A}$ is *correct* if for every instance $(G, \mathbf{x}) \in \mathcal{F}$, the resulting output vector $\mathbf{y}$ satisfies $(G, \mathbf{x}, \mathbf{y}) \in \Pi$.

Let $\mathcal{A}$ be a local deterministic algorithm for Π and $\mathcal{F}$. The *running time* of $\mathcal{A}$ over a particular instance $(G, \mathbf{x}) \in \mathcal{F}$, denoted $T_{\mathcal{A}}(G, \mathbf{x})$, is the number of rounds from the beginning of the execution of $\mathcal{A}$ until all nodes terminate. The running time of $\mathcal{A}$ is typically evaluated with respect to a collection Λ of parameters $\mathtt{q}_1, \ldots, \mathtt{q}_\ell$. Specifically, it is compared to a non-decreasing function $f\colon \mathbf{N}^\ell \to \mathbf{R}^+$; we say that f is an upper bound for the running time of $\mathcal{A}$ with respect to Λ if $T_{\mathcal{A}}(G, \mathbf{x}) \leqslant f(\mathtt{q}_1^*, \ldots, \mathtt{q}_\ell^*)$ for every instance $(G, \mathbf{x}) \in F$ with parameters $\mathtt{q}_i^* = \mathtt{q}_i(G, \mathbf{x})$ for $i \in [1, \ell]$. Let us stress that we assume throughout the paper that all the functions bounding running times of algorithms are non-decreasing.

For an integer i, the algorithm $\mathcal{A}$ *restricted to i rounds* is the local algorithm $\mathcal{B}$ that consists of running $\mathcal{A}$ for precisely i rounds. The output $\mathbf{y}(u)$ of $\mathcal{B}$ at a vertex u is defined as follows: if, during the i rounds, $\mathcal{A}$ outputs a value y at u then $\mathbf{y}(u) = y$; otherwise we let $\mathbf{y}(u)$ be an arbitrary value, e.g., "0".

A *randomized* local algorithm is a local algorithm that allows each node to use random bits in its local computation—the random bits used by different nodes being independent. A randomized (local) algorithm $\mathcal{A}$ is *Las Vegas* if its correctness is guaranteed with probability 1. The *running time* of a Las Vegas algorithm $\mathcal{A}_{LV}$ over a particular configuration $(G, \mathbf{x}) \in \mathcal{F}$, denoted $T_{\mathcal{A}_{LV}}(G, \mathbf{x})$, is a random variable, which may be unbounded. However, the expected value of $T_{\mathcal{A}_{LV}}(G, \mathbf{x})$ is bounded. A *Monte-Carlo* algorithm $\mathcal{A}_{MC}$ with guarantee $\rho \in (0\,, 1]$ is a randomized algorithm that takes a configuration $(G, \mathbf{x}) \in \mathcal{F}$ as input and terminates before a predetermined time $T_{\mathcal{A}_{MC}}(G, \mathbf{x})$ (called the *running time* of $\mathcal{A}_{MC}$). It is certain that the output vector produced by Algorithm $\mathcal{A}_{MC}$ is a solution to Π with probability at least ρ. Finally, a *weak Monte-Carlo* algorithm $\mathcal{A}_{WMC}$ with guarantee $\rho \in (0\,, 1]$ guarantees that with probability at least ρ, the algorithm outputs a correct solution by its running time $T_{\mathcal{A}_{WMC}}(G, \mathbf{x})$. (Observe that it is not certain that any execution of the weak Monte-Carlo algorithm will terminate by the prescribed time $T_{\mathcal{A}_{WMC}}(G, \mathbf{x})$, or even terminate at all.) Note that a Monte-Carlo algorithm is in particular a weak Monte-Carlo algorithm, with the same running time and guarantee. Moreover, for any constant $\rho \in (0\,, 1]$, a Las Vegas algorithm running in expected time T is a weak Monte-Carlo algorithm with guarantee ρ running in time $\frac{T}{1-\rho}$, by Markov's inequality.

Synchronicity and time complexity: Many $\mathcal{LOCAL}$ algorithms happen to have different termination times at different nodes. On the other hand, most of the algorithms rely on a simultaneous wake-up time for all nodes. This becomes an issue when one wants to run an algorithm $\mathcal{A}_1$ and subsequently an algorithm $\mathcal{A}_2$ taking the output of $\mathcal{A}_1$ as input. Indeed, this problem amounts to running $\mathcal{A}_2$ with non-simultaneous wake-up times: a node u starts $\mathcal{A}_2$ when it terminates $\mathcal{A}_1$.

As observed (e.g., by Kuhn [22]), the concept of synchronizer [2], used in the context of local algorithms, allows one to transform an asynchronous local algorithm to a synchronous one that runs in the same asymptotic time complexity. Hence, the synchronicity assumption can actually be removed. Although the standard asynchronous model introduced still assumes a simultaneous wake-up time, it can be easily verified that the technique still applies with non-simultaneous wake-up times if a node can buffer messages received before it wakes up, which is the case when running an algorithm after another.

However, we have to adapt the notion of running time. The computation that a node performs in time t depends on its interactions with nodes at distance at most t in the network. More precisely, we say that a node u terminates in time t if it terminates at most t rounds after all nodes in $B_G(u, t)$ have woken up. The termination time of u is the least t such that u terminates in time t. We finally define the running time of an algorithm as the maximum termination time over all nodes and all wake-up patterns.

Given two local algorithms $\mathcal{A}_1$ and $\mathcal{A}_2$, we let $\mathcal{A}_1; \mathcal{A}_2$ be the process of running $\mathcal{A}_2$ after $\mathcal{A}_1$. It turns out that

the running time of $\mathcal{A}_1;\mathcal{A}_2$ is bounded from above by the sum of the running times of $\mathcal{A}_1$ and $\mathcal{A}_2$. This can be shown as follows. Let t_1 and t_2 be the running times of $\mathcal{A}_1$ and $\mathcal{A}_2$ respectively. Consider a node u and let t_0 be the last wake-up time of a node in the ball $B_G(u, t_1+t_2)$. At $t_0 + t_1$, all nodes in $B_G(u, t_2)$ have terminated $\mathcal{A}_1$ and are thus considered as woken up for the execution of $\mathcal{A}_2$. Node u thus terminates before $(t_0 + t_1) + t_2$. As this is true for any node u independently of the wake-up pattern, $\mathcal{A}_1;\mathcal{A}_2$ has running time at most $t_1 + t_2$. This establishes the following observation.

Observation 2.1 *For any two local algorithms $\mathcal{A}_1$ and $\mathcal{A}_2$, the running time of $\mathcal{A}_1;\mathcal{A}_2$ is bounded by the sum of the running times of $\mathcal{A}_1$ and $\mathcal{A}_2$.*

Another useful remark is that a simultaneous wake-up algorithm running in time t can be emulated in a non-simultaneous wake-up environment with running time at most t using the simple α synchronizer. Indeed, consider a node u and let t_0 be the last wake-up time of a node in the ball $B_G(u, t)$. At time t_0, all nodes in $B_G(u, t)$ perform (or have performed) round 0. Using the α synchronizer a node can perform round i when all its neighbors have performed round $i - 1$. We can thus show by induction on i that all nodes in $B_G(u, t-i)$ perform (or have performed) round i at time t_0+i. The node u thus terminates in time t. This implies that the running time of the emulation of the algorithm with the α synchronizer is at most t. Therefore, in the remaining of the paper we may assume without loss of generality that all nodes wake up simultaneously at time 0.

Local algorithms requiring parameters: Fix a problem Π and let $\mathcal{F}$ be a collection of instances for Π. Let Γ be a collection of parameters $\mathtt{p}_1, \ldots, \mathtt{p}_r$ and let $\mathcal{A}$ be a local algorithm. We say that $\mathcal{A}$ *requires* Γ if the code of $\mathcal{A}$, which is executed by each node of the input configuration, uses a value $\tilde{\mathtt{p}}$ for each parameter $\mathtt{p} \in \Gamma$. (Note that this value is thus the same for all nodes.) The value $\tilde{\mathtt{p}}$ is a *guess* for $\mathtt{p}$. A collection of guesses for the parameters in Γ is denoted by $\tilde{\Gamma}$ and an algorithm $\mathcal{A}$ that requires Γ is denoted by $\mathcal{A}^{\Gamma}$. An algorithm that does not require any parameter is called *uniform*.

Consider an instance $(G, \mathbf{x}) \in \mathcal{F}$, a collection Γ of parameters and a parameter $\mathtt{p} \in \Gamma$. A guess $\tilde{\mathtt{p}}$ for $\mathtt{p}$ is termed *good* if $\tilde{\mathtt{p}} \geqslant \mathtt{p}(G, \mathbf{x})$, and the guess $\tilde{\mathtt{p}}$ is called *correct* if $\tilde{\mathtt{p}} = \mathtt{p}(G, \mathbf{x})$. We typically write correct guesses and collection of correct guesses with a star superscript, as in $\mathtt{p}^*$ and $\Gamma^*(G, \mathbf{x})$, respectively. When $(G, \mathbf{x})$ is clear from the context, we may use the notation Γ^* instead of $\Gamma^*(G, \mathbf{x})$.

An algorithm $\mathcal{A}^{\Gamma}$ *depends* on Γ if for every instance $(G, \mathbf{x}) \in \mathcal{F}$, the correctness of $\mathcal{A}^{\Gamma}$ over $(G, \mathbf{x})$ is guaranteed only when $\mathcal{A}^{\Gamma}$ uses a collection $\tilde{\Gamma}$ of good guesses.

Consider an algorithm $\mathcal{A}^{\Gamma}$ that depends on a collection Γ of parameters $\mathtt{p}_1, \ldots, \mathtt{p}_r$ and fix an instance $(G, \mathbf{x})$. Observe that the running time of $\mathcal{A}^{\Gamma}$ over $(G, \mathbf{x})$ may be different for different collections of guesses $\tilde{\Gamma}$, in other words, the running time over $(G, \mathbf{x})$ may be a function of $\tilde{\Gamma}$. Recall that when we consider an algorithm that does not require parameters, we still typically evaluate its running time with respect to a collection of parameters Λ. We generalize this to the case where the algorithm depends on Γ as follows.

Consider two collections Γ and Λ of parameters $\mathtt{p}_1, \ldots, \mathtt{p}_r$ and $\mathtt{q}_1, \ldots, \mathtt{q}_\ell$, respectively. Some parameters may belong to both Γ and Λ. Without loss of generality, we shall always assume that $\{\mathtt{p}_{r'+1}, \ldots, \mathtt{p}_r\} \cap \{\mathtt{q}_{r'+1}, \ldots, \mathtt{q}_\ell\} = \emptyset$ for some $r' \in [0\,, \min\{r, \ell\}]$ and $\mathtt{p}_i = \mathtt{q}_i$ for every $i \in [1\,, r']$. Notice that $\Gamma \setminus \Lambda = \{\mathtt{p}_{r'+1}, \mathtt{p}_{r'+2}, \ldots, \mathtt{p}_r\}$. A function $f\colon (\mathbf{R}^+)^\ell \to \mathbf{R}^+$ *upper bounds* the running time of $\mathcal{A}^{\Gamma}$ with respect to Γ and Λ if the running time $T_{\mathcal{A}^{\Gamma}}(G, \mathbf{x})$ of $\mathcal{A}^{\Gamma}$ for $(G, \mathbf{x}) \in \mathcal{F}$ using a collection of good guesses $\tilde{\Gamma} = \{\tilde{\mathtt{p}}_1, \ldots, \tilde{\mathtt{p}}_r\}$ is at most $f(\tilde{\mathtt{p}}_1, \ldots, \tilde{\mathtt{p}}_{r'}, \ldots, \mathtt{q}^*_\ell)$, where $\mathtt{q}^*_i = \mathtt{q}_i(G, \mathbf{x})$ for $i \in [r' + 1\,, \ell]$. Note that we do not put any restriction on the running time of $\mathcal{A}^{\Gamma}$ over $(G, \mathbf{x})$ if some of the guesses in $\tilde{\Gamma}$ are not good. In fact, in such a case, the algorithm may not even terminate and it may also produce wrong results.

For simplicity of notation, when Γ and Λ are clear from the context, we say that f upper bounds the running time of $\mathcal{A}^{\Gamma}$, without writing that it is with respect to Γ and Λ.

The set Γ is *weakly-dominated* by Λ if for each $j \in [r'+1\,, r]$, there exists an index $i_j \in [1\,, \ell]$ and an ascending function g_j such that $g_j(\mathtt{p}_j(G, \mathbf{x})) \leqslant \mathtt{q}_{i_j}(G, \mathbf{x})$ for every instance $(G, \mathbf{x}) \in \mathcal{F}$. (For example, $\Gamma = \{\Delta\}$ is weakly-dominated by $\Lambda = \{n\}$, since $\Delta(G, \mathbf{x}) \leqslant n(G, \mathbf{x})$ for any $(G, \mathbf{x})$.)

3 Pruning Algorithms

3.1 Overview

Consider a problem Π in the centralized setting and an efficient randomized Monte-Carlo algorithm $\mathcal{A}$ for Π. A known method for transforming $\mathcal{A}$ into a Las Vegas algorithm is based on repeatedly doing the following. Execute $\mathcal{A}$ and, subsequently, execute an algorithm that checks the validity of the output. If the checking fails then continue, and otherwise, terminate, i.e., break the loop. This transformation can yield a Las Vegas algorithm whose expected running time is similar to the

running time of the Monte-Carlo algorithm provided that the checking mechanism used is efficient.

If we wish to come up with a similar transformation in the context of locality, a first idea would be to consider a local algorithm that checks the validity of a tentative output vector. This concept has been studied from various perspectives (cf., e.g., [16,21,32]). However, such fast local checking procedures can only guarantee that faults are detected by at least one node, whereas to restart the Monte-Carlo algorithm, all nodes should be aware of a fault. This notification can take diameter time and will thus violate the locality constraint (i.e. running in a bounded number of rounds).

Instead of using local checking procedures, we introduce the notion of *pruning algorithms*. Informally, this is a mechanism that identifies "valid areas" where the tentative output vector $\hat{\mathbf{y}}$ is valid and *prunes* these areas, i.e., takes them out of further consideration. A pruning algorithm $\mathcal{P}$ must satisfy two properties, specifically, (1) *gluing*: $\mathcal{P}$ must make sure that the current solution on these "pruned areas" can be extended to a valid solution for the remainder of the graph, and (2) *solution detection*: if $\hat{\mathbf{y}}$ is a valid global solution to begin with then $\mathcal{P}$ should prune all nodes. Observe that since the empty output vector is a solution for the empty input graph then (1) implies the converse of (2), that is, if $\mathcal{P}$ prunes all nodes, then $\hat{\mathbf{y}}$ is a valid global solution.

Now, given a Monte-Carlo algorithm $\mathcal{A}$ and a pruning algorithm $\mathcal{P}$ for the problem, we can transform $\mathcal{A}$ into a Las Vegas algorithm by executing the pair of algorithms $(\mathcal{A};\mathcal{P})$ in iterations, where each iteration i is executed on the graph G_i induced by the set of nodes that were not pruned in previous iterations (G_1 is the initial graph G). If, in some iteration i, Algorithm $\mathcal{A}$ solves the problem on the graph G_i, then the solution detection property guarantees that the subsequent pruning algorithm will prune all nodes in G_i and hence at that time all nodes are pruned and the execution terminates. Furthermore, using induction, it can be shown that the gluing property guarantees that the correct solution to G_i combined with the outputs of the previously pruned nodes forms a solution to G.

3.2 Pruning Algorithms: Definition and Examples.

We now formally define pruning algorithms. Fix a problem Π and a family of instances $\mathcal{F}$ for Π. A *pruning* algorithm $\mathcal{P}$ for Π and $\mathcal{F}$ is a uniform algorithm that takes as input a triplet $(G,\mathbf{x},\hat{\mathbf{y}})$, where $(G,\mathbf{x})\in\mathcal{F}$ and $\hat{\mathbf{y}}$ is some tentative output vector (i.e. an output vector that may be incorrect), and returns a configuration $(G',\mathbf{x}')$ such that G' is an induced subgraph of G and $(G',\mathbf{x}')\in\mathcal{F}$. Thus, at each node v of G, the pruning algorithm $\mathcal{P}$ returns a bit $b(v)$ that indicates whether v belongs to some selected subset W of nodes of G to be pruned. (Recall that the idea is to assume that nodes in W have a satisfying tentative output value and that they can be excluded from further computations.) Note that $\mathbf{x}'$ may be different than $\mathbf{x}$ restricted to the nodes outside W.

Consider now an output vector $\mathbf{y}'$ for the nodes in $V(G')$. The *combined* output vector $\mathbf{y}$ of the vectors $\hat{\mathbf{y}}$ and $\mathbf{y}'$ is the output vector that is a combination of $\hat{\mathbf{y}}$ restricted to the nodes in W and $\mathbf{y}'$ restricted to the nodes in G', i.e., $\mathbf{y}(v)=\hat{\mathbf{y}}(v)$ if $v\in W$ and $\mathbf{y}(v)=\mathbf{y}'(v)$ otherwise. A pruning algorithm $\mathcal{P}$ for a problem Π must satisfy the following properties.

- **Solution detection:** if $(G,\mathbf{x},\hat{\mathbf{y}})\in\Pi$, then $W=V(G)$, that is, $\mathcal{P}(G,\mathbf{x},\hat{\mathbf{y}})=(\emptyset,\emptyset)$.

- **Gluing:** if $\mathcal{P}(G,\mathbf{x},\hat{\mathbf{y}})=(G',\mathbf{x}')$ and $\mathbf{y}'$ is a solution for $(G',\mathbf{x}')$, i.e., $(G',\mathbf{x}',\mathbf{y}')\in\Pi$, then the combined output vector $\mathbf{y}$ is a solution for $(G,\mathbf{x})$, i.e, $(G,\mathbf{x},\mathbf{y})\in\Pi$.

As mentioned earlier, it follows from the gluing property that if the pruning algorithm $\mathcal{P}$ returns $(\emptyset,\emptyset)$ (i.e., all nodes are pruned) then $(G,\mathbf{x},\hat{\mathbf{y}})\in\Pi$.

The pruning algorithm $\mathcal{P}$ is *monotone with respect to a parameter* $\mathtt{p}$ if $\mathtt{p}(G,\mathbf{x})\geqslant\mathtt{p}(\mathcal{P}(G,\mathbf{x},\hat{\mathbf{y}}))$ for every $(G,\mathbf{x})\in\mathcal{F}$ and every tentative output vector $\hat{\mathbf{y}}$. The pruning algorithm $\mathcal{P}$ is *monotone with respect to a collection of parameters* Γ if $\mathcal{P}$ is monotone with respect to every parameter $\mathtt{p}\in\Gamma$. In such a case, we may also say that $\mathcal{P}$ is Γ*-monotone*. The following assertions follow from the definitions.

Observation 3.1 *Let $\mathcal{P}$ be a pruning algorithm.*

1. *Algorithm $\mathcal{P}$ is monotone with respect to any non-decreasing graph-parameter.*
2. *If the configuration $(G',\mathbf{x}')$ returned by $\mathcal{P}$ satisfies $\mathbf{x}'(v)=\mathbf{x}(v)$ for every $v\in V(G)\setminus W$ and every configuration $(G,\mathbf{x})$, then $\mathcal{P}$ is monotone with respect to any non-decreasing parameter.*

For simplicity, we impose that the running time of a pruning algorithm $\mathcal{P}$ be constant. We shall elaborate on general pruning algorithms at the end of the paper.

We now give examples of pruning algorithms for several problems, namely, $(2,\beta)$-Ruling set for a constant integer β (recall that MIS is precisely $(2,1)$-Ruling set), and maximal matching. These pruning algorithms ignore the input of the nodes. Thus, by Observation 3.1, they are monotone with respect to any non-decreasing parameter.

The $(2,\beta)$-ruling set pruning algorithm: Let β be a constant integer. We define a pruning algorithm $\mathcal{P}_{(2,\beta)}$ for the $(2,\beta)$-ruling set problem as follows. Given a triplet $(G,\mathbf{x},\hat{\mathbf{y}})$, let W be the set of nodes u satisfying one of the following two conditions.

- $\hat{\mathbf{y}}(u)=1$ and $\hat{\mathbf{y}}(v)=0$ for all $v\in N(u)$, or
- $\hat{\mathbf{y}}(u)=0$ and $\exists v\in B_G(u,\beta)$ such that $\hat{\mathbf{y}}(v)=1$ and $\hat{\mathbf{y}}(w)=0$ for all $w\in N(v)$.

The question of whether a node u belongs to W can be determined by inspecting $B_G(u,1+\beta)$, the ball of radius $1+\beta$ around u. Hence, we obtain the following.

Observation 3.2 *Algorithm $\mathcal{P}_{(2,\beta)}$ is a pruning algorithm for the $(2,\beta)$-ruling set problem, running in time $1+\beta$. (In particular, $\mathcal{P}_{(2,1)}$ is a pruning algorithm for the MIS problem running in time 2.) Furthermore, $\mathcal{P}_{(2,\beta)}$ is monotone with respect to any non-decreasing parameter.*

The maximal matching problem: We define a pruning algorithm $\mathcal{P}_{\texttt{MM}}$ as follows. Given a tentative output vector $\hat{\mathbf{y}}$, recall that u and v are matched when u and v are neighbors, $\hat{\mathbf{y}}(u)=\hat{\mathbf{y}}(v)$ and $\hat{\mathbf{y}}(w)\neq\hat{\mathbf{y}}(u)$ for every $w\in(N_G(u)\cup N_G(v))\setminus\{u,v\}$. Set W to be the set of nodes u satisfying one of the following conditions.

- $\exists v\in N(u)$ such that u and v are matched, or
- $\forall v\in N(u)$, $\exists w\neq u$ such that v and w are matched.

Observation 3.3 *Algorithm $\mathcal{P}_{\texttt{MM}}$ is a pruning algorithm for MM whose running time is 3. Furthermore, $\mathcal{P}_{\texttt{MM}}$ is monotone with respect to any parameter.*

We exhibit several applications of pruning algorithms. The main application appears in the next section, where we show how pruning algorithms can be used to transform non-uniform algorithms into uniform ones. Before we continue, we need the concept of alternating algorithms.

3.3 Alternating Algorithms

A pruning algorithm can be used in conjunction with a sequence of algorithms as follows. Let $\mathcal{F}$ be a collection of instances for some problem Π. For each $i\in\mathbf{N}$, let $\mathcal{A}_i$ be an algorithm defined on $\mathcal{F}$. Algorithm $\mathcal{A}_i$ does not necessarily solve Π, it is only assumed to produce some output.

Let $\mathcal{P}$ be a pruning algorithm for Π and $\mathcal{F}$, and for $i\in\mathbf{N}$, let $\mathcal{B}_i=(\mathcal{A}_i;\mathcal{P})$, that is, given an instance $(G,\mathbf{x})$, Algorithm $\mathcal{B}_i$ first executes $\mathcal{A}_i$, which returns an output vector $\mathbf{y}$ for the nodes of G and, subsequently, Algorithm $\mathcal{P}$ is executed over the triplet $(G,\mathbf{x},\mathbf{y})$. We define the *alternating algorithm* π for $(\mathcal{A}_i)_{i\in\mathbf{N}}$ and $\mathcal{P}$ as follows. The alternating algorithm $\pi=\pi((\mathcal{A}_i)_{i\in\mathbf{N}},\mathcal{P})$ executes the algorithms $\mathcal{B}_i$ for $i=1,2,3,\ldots$ one after the other: let $(G_1,\mathbf{x}_1)=(G,\mathbf{x})$ be the initial instance given to π; for $i\in\mathbf{N}$, Algorithm $\mathcal{A}_i$ is executed on the instance $(G_i,\mathbf{x}_i)$ and returns the output vector $\mathbf{y}_i$. The subsequent pruning algorithm $\mathcal{P}$ takes the triplet $(G_i,\mathbf{x}_i,\mathbf{y}_i)$ as input and produces the instance $(G_{i+1},\mathbf{x}_{i+1})$. See Figure 1 for a schematic view of an alternating algorithm. The definition extends to a finite sequence $(\mathcal{A}_i)_{i=1}^k$ of algorithms in a natural way; the alternating algorithm for $(\mathcal{A})_{i=1}^k$ and $\mathcal{P}$ being $A_1;\mathcal{P};A_2;\mathcal{P};\cdots;A_k;\mathcal{P}$.

The alternating algorithm π *terminates* on an instance $(G,\mathbf{x})\in\mathcal{F}$ if there exists k such that $V(G_k)=\emptyset$. Observe that in such a case, the tail $\mathcal{B}_k;\mathcal{B}_{k+1};\cdots$ of π is trivial. The output vector $\mathbf{y}$ of a terminating alternating algorithm π is defined as the combination of the output vectors $\mathbf{y}_1,\mathbf{y}_2,\mathbf{y}_3,\ldots$. Specifically, for $s\in[1\,,k-1]$, let $W_s=V(G_s)\setminus V(G_{s+1})$. (Observe that W_s is precisely the set of nodes pruned by the execution of the pruning algorithm $\mathcal{P}$ in $\mathcal{B}_s$.) Then, the collection $\{W_s:1\leqslant s\leqslant k-1\}$ forms a partition of $V(G)$, i.e., $W_s\cap W_{s'}=\emptyset$ if $s\neq s'$, and $\cup_{s=1}^{k-1}W_s=V(G)$. Observe that the final output $\mathbf{y}$ of π satisfies $\mathbf{y}(u)=\mathbf{y}_s(u)$ for every node u, where s is such that $u\in W_s$. In other words, the output of π restricted to the nodes in W_s is precisely the corresponding output of Algorithm $\mathcal{A}_s$. The next observation readily follows from the definition of pruning algorithms.

Observation 3.4 *Consider a problem Π, a collection of instances $\mathcal{F}$, a sequence of algorithms $(\mathcal{A}_i)_{i\in\mathbf{N}}$ defined on $\mathcal{F}$ and a pruning algorithm $\mathcal{P}$ for Π and $\mathcal{F}$. Consider the alternating algorithm $\pi=\pi((\mathcal{A}_i)_{i\in\mathbf{N}},\mathcal{P})$ for $(\mathcal{A}_i)_{i\in\mathbf{N}}$ and $\mathcal{P}$. If π terminates on an instance $(G,\mathbf{x})\in\mathcal{F}$ then it produces a correct output $\mathbf{y}$, that is, $(G,\mathbf{x},\mathbf{y})\in\Pi$.*

In what follows, we often produce a sequence of algorithms $(\mathcal{A}_i)_{i\in\mathbf{N}}$ from an algorithm $\mathcal{A}^\Gamma$ requiring a collection Γ of non-decreasing parameters. The general idea is to design a sequence of guesses $\tilde{\Gamma}_i$ and let $\mathcal{A}_i$ be algorithm $\mathcal{A}^\Gamma$ provided with guesses $\tilde{\Gamma}_i$. Given a pruning algorithm $\mathcal{P}$, we obtain a uniform alternating algorithm $\pi=\pi((\mathcal{A}_i)_{i\in\mathbf{N}},\mathcal{P})$. The sequence of guesses is designed such that for any configuration $(G,\mathbf{x})\in\mathcal{F}$, there exists some i for which $\tilde{\Gamma}_i$ is a collection of good guesses for $(G,\mathbf{x})$. The crux is to obtain an execution time for $\mathcal{A}_1;\mathcal{P};\cdots;\mathcal{A}_i;\mathcal{P}$ of the same order as the exe-

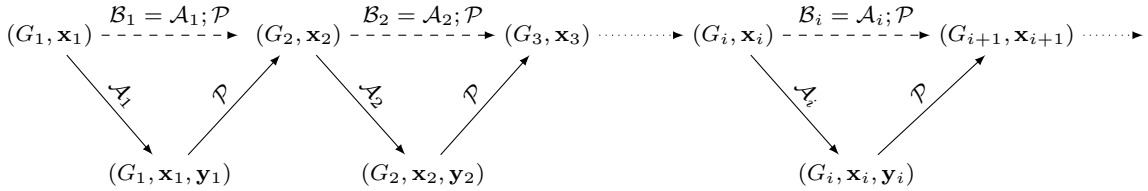

Fig. 1 Schematic view of an alternating algorithm for $(\mathcal{A}_i)_{i \in \mathbf{N}}$ and $\mathcal{P}$.

cution time of $\mathcal{A}^\Gamma$ provided with the collection $\Gamma^*(G, \mathbf{x})$ of correct guesses.

4 The General Method

We now turn to the main application of pruning algorithms discussed in this paper, that is, the construction of a transformer taking a non-uniform algorithm A^Γ as a black box and producing a uniform one that enjoys the same (asymptotic) time complexity as the original non-uniform algorithm.

We begin with a few illustrative examples of our method in Subsection 4.1. Then, the general framework of our transformer is given in Subsection 4.2. This subsection introduces a concept of "sequence-number functions" as well as the a fundamental construction used in our forthcoming algorithms.

Then, in Subsection 4.3, we consider the deterministic setting: a somewhat restrictive, yet useful, transformer is given in Theorem 1. This transformer considers a single set Γ of non-decreasing parameters $\mathtt{p}_1, \ldots, \mathtt{p}_\ell$, and assumes that (1) the given non-uniform algorithm $\mathcal{A}^\Gamma$ depends on Γ and (2) the running time of $\mathcal{A}^\Gamma$ is evaluated with respect to the parameters in Γ. Such a situation is customary, and occurs for instance for the best currently known MIS Algorithms [4,22,34] as well as for the maximal matching algorithm of Hanckowiak *et al.* [19]. As a result, the transformer given by Theorem 1 can be used to transform each of these algorithms into a uniform one with asymptotically the same time complexity.

The transformer of Theorem 1 is extended to the randomized setting in Subsection 4.4. In Subsection 4.5, we establish Theorem 3, which generalizes both Theorem 1 and Theorem 2. Finally, we conclude the section with Theorem 4 in Subsection 4.6, which shows how to manipulate several uniform algorithms that run in unknown times to obtain a uniform algorithm that runs as fast as the fastest algorithm among those given algorithms.

4.1 Some Illustrative Examples

The basic idea is very simple. Consider a problem for which we have a pruning algorithm $\mathcal{P}$, and a non uniform algorithm $\mathcal{A}$ that requires the upper bounds on some parameters to be part of the input. To obtain a uniform algorithm, we execute the pair of algorithms $(\mathcal{A}; \mathcal{P})$ in iterations, where each iteration executes $\mathcal{A}$ using a specific set of guesses for the parameters. Typically, as iterations proceed, the guesses for the parameters grow larger and larger until we reach an iteration i where all the guesses are larger than the actual value of the corresponding parameters. In this iteration, the operation of $\mathcal{A}$ on G_i using such guesses guarantees a correct solution on G_i (G_i is the graph induced by the set of nodes that were not pruned in previous iterations). The solution detection property of the pruning algorithm then guarantees that the execution terminates in this iteration and hence, Observation 3.4 guarantees that the output of all nodes combines to a global solution on G. To bound the running time, we shall make sure that the total running time is dominated by the running time of the last iteration, and that this last iteration is relatively fast.

There are various delicate points when using this general strategy. For example, in iterations where incorrect guesses are used, we have no control over the behavior of the non-uniform algorithm $\mathcal{A}$ and, in particular, it may run for too many rounds, perhaps even indefinitely. To overcome this obstacle, we allocate a prescribed number of rounds for each iteration; if Algorithm $\mathcal{A}$ reaches this time bound without outputting at some node u, then we force it to terminate with an arbitrary output. Subsequently, we run the pruning algorithm and proceed to the next iteration.

Obviously, this simple approach of running in iterations and increasing the guesses from iteration to iteration is hardly new. It was used, for example, in the context of wireless networks to compute estimates of parameters (cf., e.g., [8,31]), or to estimate the number of faults [25]. It was also used by Barenboim and Elkin [6] to avoid the necessity of having an upper bound on the arboricity a in one of their MIS algorithms, although their approach increases the running time by $\log^* n$.

One of the main contributions of the current paper is the formalization and generalization of this technique, allowing it to be used for a wide varieties of problems and applications. Interestingly, note that we are only concerned with getting rid of the use of some global parameters in the code of local algorithms, and not with obtaining estimates for them (in particular, when our algorithms terminate, a node has no guarantee to have upper bounds on these global parameters).

To illustrate the method, let us consider the non-uniform MIS algorithm of Panconesi and Srinivasan [34]. The code of Algorithm $\mathcal{A}$ uses an upper bound $\tilde{n}$ on the number of nodes n, and runs in time at most $f(\tilde{n}) = 2^{O(\sqrt{\log \tilde{n}})}$. Consider a pruning algorithm $\mathcal{P}_{\texttt{MIS}}$ for MIS (such an algorithm is given by Observation 3.2). The following sketches our technique for obtaining a uniform MIS algorithm. For each integer i, set $n_i = \max\left\{a \in \mathbf{N} : f(a) \leqslant 2^i\right\}$. In Iteration i, for $i = 1, 2, \ldots$, we first execute Algorithm $\mathcal{A}$ using the guess n_i (as an input serving as an upper bound for the number of nodes) for precisely 2^i rounds. Subsequently, we run the pruning algorithm $\mathcal{P}_{\texttt{MIS}}$. When the pruning algorithm terminates, we execute the next iteration on the non-pruned nodes. Let s be the integer such that $2^{s-1} < f(n) \leqslant 2^s$, where n is the number of nodes of the input graph. By the definition, $n \leqslant n_s$. Therefore, the application of $\mathcal{A}$ in Iteration s uses a guess n_s that is indeed good, i.e., larger than the number of nodes. Moreover, this execution of $\mathcal{A}$ is completed before the prescribed deadline of 2^s rounds expires because its running time is at most $f(n_s) \leqslant 2^s$. Hence, we are guaranteed to have a correct solution by the end of Iteration s. The running time is thus at most $\sum_{i=1}^{s} 2^i = O(f(n))$.

This method can sometimes be extended to simultaneously remove the use of several parameters in the code of a local algorithm. For example, consider the MIS algorithm of Barenboim and Elkin [4] (or that of Kuhn [22]), which uses upper bounds $\tilde{n}$ and $\tilde{\Delta}$ on n and Δ, respectively, and runs in time $f(\tilde{n}, \tilde{\Delta}) = f_1(\tilde{n}) + f_2(\tilde{\Delta})$, where $f_1(\tilde{\Delta}) = O(\tilde{\Delta})$ and $f_2(\tilde{n}) = O(\log^* \tilde{n})$. The following sketches our method for obtaining a corresponding uniform MIS algorithm that runs in time $O(f(n, \Delta))$. For each integer i, set $n_i = \max\left\{a \in \mathbf{N} : f_1(a) \leqslant 2^i\right\}$ and $\Delta_i = \max\left\{a \in \mathbf{N} : f_2(a) \leqslant 2^i\right\}$. In Iteration i, for $i = 1, 2, \ldots$, we first execute Algorithm $\mathcal{A}$ using the guesses n_i and Δ_i, but this time the execution lasts for precisely $2 \cdot 2^i$ rounds. (The factor 2 in the running time of an iteration follows from the fact that the running time is the sum of two non-negative ascending functions of two different parameters, namely $f_1(n)$ and $f_2(\Delta)$.) Subsequently, we run the pruning algorithm $\mathcal{P}_{\texttt{MIS}}$, and as before, when the pruning algorithm terminates, we execute the next iteration on the non-pruned nodes. Now, let s be the integer such that $2^{s-1} < f(n, \Delta) \leqslant 2^s$. By the definition, $n \leqslant n_s$ and $\Delta \leqslant \Delta_s$. Hence, the application of $\mathcal{A}$ in Iteration s uses guesses that are indeed good. This execution of A is completed before the prescribed deadline of 2^{s+1} rounds expires because its running time is at most $f_1(n_s) + f_2(\Delta_s) \leqslant 2^{s+1}$. Thus, the algorithm consists of at most s iterations. Since the running time of the whole execution is dominated by the running time of the last iteration, the total running time is $O(2^{s+1}) = O(f(n, \Delta))$.

The above discussion shall be formalized in Theorem 1. Before stating and proving it, though, we need one more concept, called "sequence-number function", which gives a certain measure for the "separation" between the variables in a function defined over $\mathbf{N}^\ell$.

4.2 The General Framework

Consider a function $f\colon \mathbf{N}^\ell \to \mathbf{R}^+$. A *set-sequence for* f is a sequence $(S_f(i))_{i \in \mathbf{N}}$ such that for every $i \in \mathbf{N}$,

(i) $S_f(i)$ is a finite subset (possibly empty) of $\mathbf{N}^\ell$; and
(ii) if $\underline{y} \in \mathbf{N}^\ell$ and $f(\underline{y}) \leqslant i$, then $\underline{y}$ is dominated by a vector $\underline{x}$ that belongs to $S_f(i)$.

The set-sequence $(S_f(i))_{i \in \mathbf{N}}$ is *bounded* if there exists a positive number c such that

$$\forall i \in \mathbf{N}, \forall \underline{x} \in S_f(i), \quad f(\underline{x}) \leqslant c \cdot i.$$

The constant c is referred to as the *bounding constant of* $(S_f(i))_{i \in \mathbf{N}}$. Note that a set-sequence may contain empty sets.

A function $s_f\colon \mathbf{N} \to \mathbf{N}$ is a *sequence-number function for* f if

(1) s_f is moderately-slow; and
(2) there exists a bounded set-sequence $(S_f(i))_{i \in \mathbf{N}}$ for f such that

$$\forall i \in \mathbf{N}, \quad |S_f(i)| \leqslant s_f(i).$$

For example, consider the case where $f\colon \mathbf{N}^\ell \to \mathbf{R}$ is additive, i.e., $f(x_1, \ldots, x_\ell) = \sum_{k=1}^{\ell} f_k(x_k)$, where $f_1, \ldots, f_\ell$ are non-negative ascending functions. Here, the constant function 1 is a sequence-number function for f. Indeed, for $i \in \mathbf{N}$, let $S_f(i) = \{\underline{x}\}$, where the k-th coordinate of $\underline{x}$ is defined to be the largest integer y such that $f_k(y) \leqslant i$ (if such an integer y exists, otherwise, $S_f(i)$ is empty). Hence, if $f(\underline{y}) \leqslant i$ then we deduce that $f_k(\underline{y}_k) \leqslant i$ as each of the functions $f_1, \ldots, f_\ell$ is non-negative. Therefore, $\underline{x}$ dominates $\underline{y}$. Consequently, $(S_f(i))_{i \in \mathbf{N}}$ is a set-sequence for f, which is bounded since

$$f(\underline{x}) \leqslant \sum_{k=1}^{\ell} f_k(\underline{x}_k) \leqslant \ell \cdot i,$$

and ℓ does not depend on i (the bounding constant c is equal to ℓ in this case).

As another example, consider the case where $f\colon \mathbf{N}^2 \to \mathbf{R}$ is given by $f(x_1, x_2) = f_1(x_1) \cdot f_2(x_2)$, where f_1 and f_2 are ascending functions taking values at least 1. Then, the function $s_f(i) = \lceil \log i \rceil + 1$ is a sequence-number for f. Indeed, for $i \in \mathbf{N}$ let $S_f(i) = \left\{(x_1^j, x_2^j) \,:\, j \in [0\,,\lceil \log i \rceil]\right\}$ where x_1^j is the largest integer y_1 such that $f_1(y_1) \leqslant 2^j$ and x_2^j is the largest integer y_2 such that $f_2(y_2) \leqslant 2^{\log i - j + 1}$ for each $j \in [0\,,\lceil \log i \rceil]$ (if such integers y_1 and y_2 exist, otherwise we do not define the pair (x_1^j, x_2^j)). Again, a straightforward check ensures that $(S_f(i))_{i \in \mathbf{N}}$ is a bounded set-sequence for f with bounding constant 2. On the other hand, it is interesting to note that not all functions have a bounded sequence-number function, as one can see by considering the min function over $\mathbf{N}^2$. The following observation summarizes to two aforementioned examples.

Observation 4.1

- *The constant function* 1 *is a sequence-number function for any additive function.*
- *Let* $f\colon \mathbf{N}^2 \to \mathbf{R}$ *be a function given by* $f(x_1, x_2) = f_1(x_1) \cdot f_2(x_2)$*, where* $f_1 \geqslant 1$ *and* $f_2 \geqslant 1$ *are ascending functions. Then, the function* $s_f(i) = \lceil \log i \rceil + 1$ *is a sequence-number function for* f*.*

We now give an explicit construction of a local algorithm π, which will be used to prove the forthcoming theorems.

Consider a problem Π and a family of instances $\mathcal{F}$. Assume that $\mathcal{P}$ is a pruning algorithm for Π. Let $\mathcal{A}^\Gamma$ be a deterministic algorithm for Π and $\mathcal{F}$ depending on a set Γ of parameters $\mathtt{p}_1, \ldots, \mathtt{p}_\ell$. In addition, fix an integer c and let $(S_i)_{i \in \mathbf{N}}$ be a family of (possibly empty) subsets of $\mathbf{N}^\ell$.

The algorithm π runs in iterations, each of which can be seen as a uniform alternating algorithm that operates on the configurations in $\mathcal{F}$.

Fix $i \in \mathbf{N}$ and let us write $S_i = \{\underline{x}^1, \ldots, \underline{x}^{J_i}\}$. For every $j \in [1\,, J_i]$, consider the uniform algorithm $\mathcal{A}_{j,i}$ that consists of running $\mathcal{A}^\Gamma$ with the vector of guesses $\underline{x}^j$ of S_i. More precisely, the k-th coordinate of $\underline{x}^j$ is used as a guess for $\mathtt{p}_k$ for $k \in \{1, \ldots, \ell\}$. Now, we define $\mathcal{A}'_{j,i}$ to be the algorithm $\mathcal{A}_{j,i}$ restricted to $c \cdot 2^i$ rounds.

An iteration of π consists of running the uniform alternating algorithm for the sequence of uniform algorithms $\{\mathcal{A}'_{j,i}\}_{j \in [1, J_i]}$ and the pruning algorithm $\mathcal{P}$. A pseudocode description of Algorithm π is given by Algorithm 1.

```
begin
    (S_f(i))_{i∈N} ⟵ bounded set-sequence for f
    corresponding to s_f;
    c ⟵ bounding constant of (S_f(i))_{i∈N};
    (G_1, x_1) ⟵ (G, x);
    for i from 1 to ∞ do
        S_i ⟵ S_f(2^i);
        J_i ⟵ |S_i|;
        (G_{1,i}, x_{1,i}) ⟵ (G_i, x_i);
        for j from 1 to J_i do
            A'_{j,i} ⟵ A^Γ restricted to c · 2^i rounds run
            with vector guesses x^j of S_i;
            y_{j,i} ⟵ A'_{j,i}(G_{j,i}, x_{j,i});
            (G_{j+1,i}, x_{j+1,i}) ⟵ P(G_{j,i}, x_{j,i}, y_{j,i});
        end
        (G_{i+1}, x_{i+1}) ⟵ (G_{J_i+1,i}, x_{J_i+1,i});
    end
end
```

Algorithm 1: The algorithm π.

We are now ready to state and prove Theorem 1, which deals with deterministic local algorithms.

4.3 The Deterministic Case

Theorem 1 considers a single set Γ of non-decreasing parameters $\mathtt{p}_1, \ldots, \mathtt{p}_\ell$, and assumes that (1) the given non-uniform algorithm $\mathcal{A}^\Gamma$ depends on Γ and (2) the running time of $\mathcal{A}^\Gamma$ is evaluated according to the parameters in Γ. Recall that in such a case, we say that a function $f\colon \mathbf{N}^\ell \to \mathbf{R}^+$ *upper bounds* the running time of $\mathcal{A}^\Gamma$ with respect to Γ if the running time $T_{\mathcal{A}^\Gamma}(G, \mathbf{x})$ of $\mathcal{A}^\Gamma$ for every $(G, \mathbf{x}) \in \mathcal{F}$ using a collection of good guesses $\tilde{\Gamma} = \{\tilde{\mathtt{p}}_1, \ldots, \tilde{\mathtt{p}}_\ell\}$ for $(G, \mathbf{x})$ is at most $f(\tilde{\mathtt{p}}_1, \ldots, \tilde{\mathtt{p}}_\ell)$.

Theorem 1 *Consider a problem* Π *and a family of instances* $\mathcal{F}$*. Let* $\mathcal{A}^\Gamma$ *be a deterministic algorithm for* Π *and* $\mathcal{F}$ *depending on a set* Γ *of non-decreasing parameters. Suppose that the running time of* $\mathcal{A}^\Gamma$ *is bounded from above by some function* $f\colon \mathbf{N}^\ell \to \mathbf{R}^+$ *where* $\ell = |\Gamma|$*. Assume that there exists a sequence-number function* s_f *for* f*, and a* Γ*-monotone pruning algorithm* $\mathcal{P}$ *for* Π *and* $\mathcal{F}$*. Then there exists a* uniform *deterministic algorithm for* Π *and* $\mathcal{F}$ *whose running time is* $O(f^* \cdot s_f(f^*))$*, where* $f^* = f(\Gamma^*)$*.*

Proof Let $\mathtt{p}_1, \ldots, \mathtt{p}_\ell$ be the parameters in Γ. Fix a bounded set-sequence $(S_f(i))_{i \in \mathbf{N}}$ for f corresponding to s_f and let c be the bounding constant of $(S_f(i))_{i \in \mathbf{N}}$. Set $S_i = S_f(2^i)$ and $J_i = |S_i|$, hence $J_i \leqslant s_f(2^i)$.

The desired uniform algorithm is the algorithm π (Algorithm 1). We shall prove that π is correct and runs in time $O(s_f(2^m) \cdot 2^m)$ over every configuration in $\mathcal{F}$, where $m = \lceil \log f^* \rceil$.

Fix $i \in \mathbf{N}$ and let us write $S_i = \{\underline{x}^1, \ldots, \underline{x}^{J_i}\}$. Each iteration of the inner loop of π is called *Sub-iteration*,

while *Iteration* is reserved for iterations of the outer loop. As written in the pseudocode description of π given by Algorithm 1, $(G_{j,i}, \mathbf{x}_{j,i})$ is the configuration over which π operates during Sub-iteration j of Iteration s, for $j \in [1\,,J_i]$.

Let us prove that Algorithm π is correct. Fix a configuration $(G, \mathbf{x})$ and set $\mathtt{p}_r^* = \mathtt{p}_r(G, \mathbf{x})$ for $r \in [1\,,\ell]$. We consider the operation of π on $(G, \mathbf{x})$. Setting $f^* = f(\mathtt{p}_1^*, \ldots, \mathtt{p}_\ell^*)$, we know that f^* is an upper bound on the running time of $\mathcal{A}^\Gamma$ over $(G, \mathbf{x})$, assuming that $\mathcal{A}^\Gamma$ uses the vector Γ^* of correct guesses $\mathtt{p}_1^*, \ldots, \mathtt{p}_\ell^*$. Let s be the least integer such that $f^* \leqslant 2^s$. By the definition, there exists $j^* \in [1\,,J_s]$, such that $\underline{x}^{j^*}$ dominates $(\mathtt{p}_1^*, \ldots, \mathtt{p}_\ell^*)$.

The monotonicity property of $\mathcal{P}$ implies that $\mathtt{p}_r(G_{j-1,i}, \mathbf{x}_{j-1,i}) \geqslant \mathtt{p}_r(G_{j,i}, \mathbf{x}_{j,i})$ for every $r \in [1\,,\ell]$. Thus, we infer by induction on k that $\mathtt{p}_r^* = \mathtt{p}_r(G, \mathbf{x}) \geqslant \mathtt{p}_r(G_{j,i}, \mathbf{x}_{j,i})$ for every $i \in \mathbf{N}$, $j \in [1\,,J_i]$ and $r \in [1\,,\ell]$.

Now, let us consider Iteration s of π. Assume that some nodes are still active during Iteration s of π, that is, $V(G_s)$ is not empty. Iteration s of π is composed of J_s sub-iterations. During Sub-iteration j, the algorithm $\mathcal{A}'_{j,s}; \mathcal{P}$ is executed over $(G_s^j, \mathbf{x}_s^j)$. We know that $\mathtt{p}_r^* \geqslant \mathtt{p}_r(G_{j,s}, \mathbf{x}_{j,s})$ for every $j \in [1, J_s]$, and every $r \in [1\,,\ell]$. So, in Sub-iteration j^* of Iteration s, we have $x_{j^*,r} \geqslant \mathtt{p}_r^* \geqslant \mathtt{p}_r(G_{j^*,s}, \mathbf{x}_{j^*,s})$ for every $r \in [1\,,\ell]$.

Sub-iteration j^* consists of first running Algorithm $\mathcal{A}'_{j^*,s}$, which amounts to running $\mathcal{A}^\Gamma$ for $c \cdot 2^s$ rounds using the vector of guesses $\underline{x}^{j^*}$ By the definition of $S_f(2^s)$, it follows that $f(\underline{x}^{j^*}) \leqslant c \cdot 2^s$. Hence, this execution of Algorithm $\mathcal{A}^\Gamma$ is actually completed by time $c \cdot 2^s$. Furthermore, since $\underline{x}^{j^*}$ dominates $(\mathtt{p}_1(G_{j^*,s}, \mathbf{x}_{j^*,s}), \ldots, \mathtt{p}_\ell(G_{j^*,s}, \mathbf{x}_{j^*,s}))$, the vector of guesses used by Algorithm $\mathcal{A}^\Gamma$ is good, and hence the algorithm outputs a vector $\mathbf{y}_s^{j^*}$ satisfying $(G_{j^*,s}, \mathbf{x}_{j^*,s}, \mathbf{y}_{j^*,s}) \in \Pi$. By the solution detection property, the subsequent pruning algorithm (still in Sub-iteration j^* of Iteration s) selects $W_{j^*,s} = V(G_{j^*,s})$. By Observation 3.4, it follows that π is correct.

It remains to prove that the running time is $O(s_f(f^*) \cdot f^*)$. Let T_0 be the running time of $\mathcal{P}$. Observe that Iteration i of π takes at most $J_i(c \cdot 2^i + T_0)$ rounds, which is $O(s_f(2^i) \cdot 2^i)$ rounds. Since π consists of at most s iterations, the running time of π is bounded by $\sum_{i=1}^{s} s_f(2^i) \cdot 2^i$, which is $O(s_f(2^s) \cdot 2^s)$ because s_f is non-decreasing. Moreover,

$$O(s_f(2^s) \cdot 2^s) = O(s_f(2 \cdot f^*) \cdot 2^s) = O(s_f(f^*) \cdot f^*)$$

since $2^{s-1} < f^* \leqslant 2^s$ and s_f is moderately-slow (hence, in particular, non-decreasing). Therefore, the running time of π is bounded by $O(s_f(f^*) \cdot f^*)$. □

By Observation 4.1, the constant function $s_f = 1$ is a sequence number function for any additive function f. Hence, Corollary 1(vi) follows directly by applying Theorem 1 to the maximal matching algorithm of Hanckowiak *et al.* [19], and using Observation 3.3.

In addition, using Observation 3.2, Theorem 1 allows us to transform each of the MIS algorithms in [4, 22,34] into a uniform one with asymptotically the same time complexity. We thus obtain the following corollary.

Corollary 2 *Consider the family $\mathcal{F}$ of all graphs.*

- *There exists a uniform deterministic* `MIS` *algorithm for $\mathcal{F}$ running in time $O(\Delta + \log^* n)$.*
- *There exists a uniform deterministic* `MIS` *algorithm for $\mathcal{F}$ running in time $2^{O(\sqrt{\log n})}$.*

Recall that Barenboim and Elkin [5] devised, for every $\delta > 0$, a (non-uniform) deterministic `MIS` algorithm for the family of all graphs running in time $f(a, n) = O(a + a^\delta \log n)$. Fix $\epsilon \in (0, 1)$ and consider the family F_{large} of graphs with arboricity $a > \log^{1+\epsilon/2} n$. It follows from [5] (applied with, e.g., $\delta = \epsilon/3$), that there exists a (non-uniform) deterministic `MIS` algorithm for F_{large} running in time $O(a)$. Hence, using Observation 3.2 and Theorem 1, we obtain a uniform deterministic `MIS` algorithm for F_{large} running in $O(a)$ time.

Next, let F_{med} be the family of graphs with arboricity a such that $\log^{1/3} n < a \leqslant \log^{1+\epsilon/2} n$. Since $a \leqslant \log^{1+\epsilon/2} n$, it follows that $a^{1-\epsilon/2} < \log n$, and hence, $a < a^{\epsilon/2} \log n$. By [5], applied with $\delta = \epsilon/2$, there exists a deterministic `MIS` algorithm for F_{med} running in time $f_{\text{med}} = O(a^{\epsilon/2} \log n)$. Note that by Observation 4.1, the sequence number for f_{med} is $s_{f_{\text{med}}}(f_{\text{med}}) = O(\log f_{\text{med}}) = O(\log \log n)$. Hence, by combining Observation 3.2 and Theorem 1, we obtain a uniform `MIS` algorithm for F_{med} running in time $O(a^{\epsilon/2} \log n \log \log n) = O(a^\epsilon \log n)$. (This last equality follows from the fact that $\log^{1/3} n < a$.)[1] Summarizing the above discussion, we obtain the following.

Corollary 3 *For every $\epsilon > 0$, there exists the following uniform deterministic* `MIS` *algorithm:*

- *For the family F_{large}, running in $O(a)$ time,*
- *For the family F_{med}, running in $O(a^\epsilon \log n)$ time.*

4.4 The Randomized Case

We now show how to extend Theorem 1 to the randomized setting. More specifically, we replace the given

[1] In fact, we could have used in the definition of F_{med} any small constant instead of 1/3, but 1/3 is sufficiently good for our purposes as, anyway, this result will be combined with better results for $a = o(\sqrt{\log n})$, which shall be established later on, in Corollary 4.

non-uniform deterministic algorithm of Theorem 1 by a non-uniform weak Monte-Carlo algorithm $\mathcal{A}^\Gamma$ and produce a uniform Las Vegas one. This transformer is more sophisticated than the one given in Theorem 1, and requires the use of sub-iterations for bounding the expected running time and probability of success of the resulting Las-Vegas algorithm.

Theorem 2 *Consider a problem Π and a family of instances $\mathcal{F}$. Let $\mathcal{A}^\Gamma$ be a weak Monte-Carlo algorithm for Π and $\mathcal{F}$ depending on a set Γ of non-decreasing parameters. Suppose that the running time of $\mathcal{A}^\Gamma$ is bounded from above by some function $f\colon \mathbf{N}^\ell \to \mathbf{R}^+$, where $\ell = |\Gamma|$. Assume that there exists a sequence-number function s_f for f, and a Γ-monotone pruning algorithm $\mathcal{P}$ for Π and $\mathcal{F}$. Then there exists a* uniform *Las Vegas algorithm for Π and $\mathcal{F}$ whose expected running time is $O(f^* \cdot s_f(f^*))$, where $f^* = f(\Gamma^*)$.*

Proof Let $\mathtt{p}_1, \ldots, \mathtt{p}_\ell$ be the parameters in Γ. Let T_0 be the running time of the pruning algorithm $\mathcal{P}$, and let $\mathcal{A}^\Gamma$ be the given weak Monte-Carlo algorithm. To simplify the notations, we assume that the success guarantee ρ of $\mathcal{A}^\Gamma$ is 1/2.

```
begin
    (S_f(i))_{i∈N} ⟵ bounded set-sequence for f
    corresponding to s_f;
    c ⟵ bounding constant of (S_f(i))_{i∈N};
    (G_1, x_1) ⟵ (G, x);
    for i from 1 to ∞ do
        for j from 1 to i do
            S_j ⟵ S_f(2^j);
            J_j ⟵ |S_j|;
            (G_{1,j}, x_{1,j}) ⟵ (G_i, x_i);
            for k from 1 to J_j do
                A'_{k,j} ⟵ A^Γ restricted to c·2^j rounds
                run with vector guesses x^k of S_j;
                y_{k,j} ⟵ A'_{k,j}(G_{k,j}, x_{k,j});
                (G_{k+1,j}, x_{k+1,j}) ⟵
                P(G_{k,j}, x_{k,j}, y_{k,j});
            end
            (G_{j+1}, x_{j+1}) ⟵ (G_{J_j+1,j}, x_{J_j+1,j});
        end
    end
end
```

Algorithm 2: The algorithm τ in the proof of Theorem 2.

The desired uniform algorithm τ runs in iterations, where Iteration i consists of running the first i iterations of the algorithm π defined in Subsection 4.2. A pseudocode description of Algorithm τ is given by Algorithm 2. Similarly as in the proof of Theorem 1, the word "Iteration" is reserved for the iterations of the outer loop of τ, while "Sub-iteration" is used for the iterations of the middle loop of τ.

For each positive integer i, let β_i be the number of rounds used in Iteration i of τ. Analogously to the proof of Theorem 1, we infer that $\beta_i = O(s_f(2^i) \cdot 2^i)$. Let α_i be the number of rounds used during the first i iterations of τ. We thus have $\alpha_i = \sum_{k=1}^{i} \beta_k$, which is $O(s_f(2^i) \cdot 2^i)$.

It follows using similar arguments to the ones given in the proof of Theorem 1, that if τ outputs, then the output vector $\mathbf{y}$ is a solution, i.e. $(G, \mathbf{x}, \mathbf{y}) \in \Pi$.

It remains to bound the running time of τ. We consider the random variable $\mathrm{T}_\tau(G, x)$ that stands for "the running time of τ on $(G, \mathbf{x})$". For every integer i, let ρ_i be the probability that $V(G_i) \neq \emptyset$ and $V(G_{i+1}) = \emptyset$, that is, ρ_i is the probability that the last active node becomes inactive precisely during Iteration i of τ. In other words,

$$\rho_i = \mathbf{Pr}\left(\mathrm{T}_\tau(G, x) \in [\alpha_{i-1} + 1\,, \alpha_i]\right).$$

Setting $f^* = f(\mathtt{p}_1^*, \ldots, \mathtt{p}_\ell^*)$, we know that f^* is an upper bound on the running time of $\mathcal{A}^\Gamma$ over $(G, \mathbf{x})$, assuming that $\mathcal{A}^\Gamma$ uses the collection Γ^* of correct guesses $\mathtt{p}_1^*, \ldots, \mathtt{p}_\ell^*$. Consider the smallest integer s such that $f^* \leqslant 2^s$.

Since s_f is moderately-slow, there is a constant K such that $\alpha_{i+1} \leqslant K \cdot \alpha_i$ for every positive integer i. In particular, $\alpha_{s+i} \leqslant K^i \cdot \alpha_s$, and hence

$$\begin{aligned}\mathbf{E}(\mathrm{T}_\tau(G, x)) &\leqslant \alpha_s \cdot \mathbf{Pr}\left(\mathrm{T}_\tau(G, \mathbf{x}) \leqslant \alpha_s\right) + \sum_{i=1}^{\infty} \alpha_{s+i} \cdot \rho_{s+i} \\ &\leqslant \alpha_s + \alpha_s \sum_{i=1}^{\infty} K^i \cdot \rho_{s+i}.\end{aligned}$$

Our next goal is to bound ρ_{s+i} from above. For a positive integer r, let χ_r be the event that $V(G_{r+1}) \neq \emptyset$, that is, none of $\mathcal{C}_1, \ldots, \mathcal{C}_r$ output the empty configuration and thus, there is still an active node at the beginning Iteration $r+1$ of τ. Thus, $\rho_{s+i} \leqslant \mathbf{Pr}(\chi_{s+i-1})$.

Recall that we assume that the success guarantee of $\mathcal{A}^\Gamma$ is 1/2. Therefore, using similar analysis as in the proof of Theorem 1, it follows that for every positive integer k, the probability that an application of $\mathcal{B}_{s+k-1}$ (in particular, during iteration $s+i-1$) does not output the empty configuration is at most 1/2. As a result,

$$\rho_{s+i} \leqslant \mathbf{Pr}(\chi_{s+i-1}) \leqslant \prod_{j=1}^{i} 2^{-j} = 2^{-(i^2+i)/2}.$$

Therefore,

$$\begin{aligned}\mathbf{E}(\mathrm{T}_\tau(G, x)) &\leqslant \alpha_s \left(1 + \sum_{i=1}^{\infty} K^i \cdot 2^{-(i^2+i)/2}\right) \\ &= O(\alpha_s) = O(f^* \cdot s_f(f^*)).\end{aligned}$$

□

Corollary 1(vii) follows by applying Theorem 2 to the ruling set algorithm of Schneider and Wattenhofer [36], and using the pruning algorithm given by Observation 3.2.

4.5 The General Theorem

Some complications arise when the correctness of the given non-uniform algorithm relies on the use of a set of parameters Γ while its running time is evaluated with respect to another set of parameters Λ. For example, it may be the case that an upper bound on a parameter $\mathtt{p}$ is required for the correct operation of an algorithm, yet the running time of the algorithm does not depend on $\mathtt{p}$. In this case, it may not be clear how to choose the guesses for $\mathtt{p}$. (This occurs, for example, in the MIS algorithms of Barenboim and Elkin [6], where the knowledge of n and the arboricity a are required, yet the running time f is a function of n only.) Such complications can be solved when there is some relation between the parameters in Γ and those in Λ; specifically, when Γ is weakly-dominated by Λ. (The definition of weakly-dominated is given in Section 2.) This issue is handled in the following theorem, which extends both Theorem 1 and Theorem 2.

Theorem 3 *Consider a problem Π, a family of instances $\mathcal{F}$ and two sets of non-decreasing parameters Γ and Λ, where Γ is weakly-dominated by Λ. Let $\mathcal{A}^\Gamma$ be a deterministic (respectively, weak Monte-Carlo) algorithm depending on Γ whose running time is upper bounded by some function $f\colon \mathbf{N}^\ell \to \mathbf{R}^+$, where $\ell = |\Lambda|$. Assume that there exists a sequence-number function s_f for f, and a $\Lambda\cup\Gamma$-monotone pruning algorithm $\mathcal{P}$ for Π and $\mathcal{F}$. Then there exists a* uniform *deterministic (resp., Las Vegas) algorithm for Π and $\mathcal{F}$ whose running time on every configuration $(G,\mathbf{x}) \in \mathcal{F}$ is $O(f^* \cdot s_f(f^*))$, where $f^* = f(\Lambda^*(G,\mathbf{x}))$.*

Proof First, we consider the case where $\Gamma \subseteq \Lambda$ and next the general case.

Assume that $\Lambda = \{\mathtt{p}_1, \ldots, \mathtt{p}_\ell\}$ and $\Gamma = \{\mathtt{p}_1, \ldots, \mathtt{p}_r\}$, where $r \leqslant \ell$. Then, let us simply impose that A^Γ also requires estimates for the parameters $\mathtt{p}_{r+1}, \ldots, \mathtt{p}_\ell$, that is, the operation of A^Γ requires such estimates but actually ignores them after obtaining them. This way, we obtain an algorithm A^Λ depending on Λ. Since f is non-decreasing, $f(\mathtt{p}_1^*, \ldots, \mathtt{p}_\ell^*) \leqslant f(\mathtt{p}_1^*, \ldots, \mathtt{p}_r^*, \tilde{\mathtt{p}}_{r+1}, \ldots, \tilde{\mathtt{p}}_\ell)$, where $\tilde{\mathtt{p}}_i$ is a good guess for every $i \in [r+1\,,\ell]$. Hence, the running time of Algorithm A^Λ is also bounded by f, so the conclusion follows by applying Theorems 1 and 2.

Now, let $\mathtt{p}_1, \ldots, \mathtt{p}_r$ and $\mathtt{q}_1, \ldots, \mathtt{q}_\ell$ be the parameters in Γ and Λ, respectively. Recall that $r' \in [0\,, \min\{r, \ell\}]$ is such that $\{\mathtt{p}_{r'+1}, \mathtt{p}_{r'+2}, \ldots, \mathtt{p}_r\} \cap \{\mathtt{q}_{r'+1}, \mathtt{q}_{r'+2}, \ldots, \mathtt{q}_\ell\} = \emptyset$ and $\mathtt{p}_i = \mathtt{q}_i$ for every $i \in [1\,, r']$. Set $t = r - r'$. As Γ is weakly-dominated by Λ, there exists a function $h : [1\,, t] \to [1\,, \ell]$ and, for each $j \in [1\,, t]$, an ascending function g_j such that $g_j(\mathtt{p}_{r'+j}(G,\mathbf{x})) \leqslant \mathtt{q}_{h(j)}(G,\mathbf{x})$ for every configuration $(G,\mathbf{x}) \in \mathcal{F}$. For every real number x, we set $g_j^{-1}(x) = \min g_j^{-1}(\{x\})$. Since g_j is ascending, $g_j^{-1}(x) \geqslant g_j^{-1}(y)$ whenever $x \geqslant y$.

Let $\Lambda' = \Lambda \cup \Gamma = \{\mathtt{q}_1, \ldots, \mathtt{q}_\ell, \mathtt{p}_{r'+1}, \ldots, \mathtt{p}_r\}$, and recall that $f\colon \mathbf{N}^\ell \to \mathbf{R}^+$ is the (non-decreasing) function bounding the running time of $\mathcal{A}^\Gamma$. We define a new function $f' : \mathbf{N}^{\ell+t} \to \mathbf{R}$ by setting

$$f'(x_1, \ldots, x_\ell, y_1, \ldots, y_t) = f(z_1, \ldots, z_\ell),$$

where for each $i \in [1\,, \ell]$,

$$z_i = \max\left(\{x_i\} \cup \left\{g_k(y_k) \,:\, k \in h^{-1}(\{i\})\right\}\right).$$

Let s_f be a sequence-number function for f and let $(S_f(i))_{i\in\mathbf{N}}$ be a corresponding bounded set-sequence with bounding constant c.

We assert that s_f is also a sequence-number function of f' and admits a corresponding bounded set-sequence with bounding constant c. To see this, we first define for $i \in \mathbf{N}$ a set $S_{f'}(i)$ with $|S_{f'}(i)| = |S_f(i)|$ as follows. For each $(x_1, \ldots, x_\ell) \in S_f(i)$, let $S_{f'}(i)$ contain $(x_1, \ldots, x_\ell, y_1, \ldots, y_t)$, where $y_j = g_j^{-1}(x_{h(j)})$ for $j \in [1\,, t]$. Observe that $g_j(y_j) = x_{h(j)}$ for every $j \in [1\,, t]$. Hence, $f'(x_1, \ldots, x_\ell, y_1, \ldots, y_t) = f(x_1, \ldots, x_\ell)$ if $(x_1, \ldots, x_\ell, y_1, \ldots, y_t) \in S_{f'}(i)$.

This observation directly implies that $f'(\underline{x}') \leqslant c \cdot i$ if $\underline{x}' \in S_{f'}(i)$, since $f(\underline{x}) \leqslant c \cdot i$ if $\underline{x} \in S_f(i)$. Now, assume that $f'(\underline{x}) \leqslant i$ for some $\underline{x} = (x_1, \ldots, x_\ell, y_1, \ldots, y_\ell) \in \mathbf{N}^{\ell+t}$. Then, $f(z_1, \ldots, z_\ell) \leqslant i$, where z_i is given by the definition of f'. Consequently, there exists a vector $\underline{\tilde{z}} \in S_f(i)$ that dominates $(z_1, \ldots, z_\ell)$. Moreover,

$$\underline{\tilde{z}}' = (\underline{\tilde{z}}_1, \ldots, \underline{\tilde{z}}_\ell, g_1^{-1}(\underline{\tilde{z}}_{h(1)}), \ldots, g_t^{-1}(\underline{\tilde{z}}_{h(t)})) \in S_{f'}(i).$$

Therefore, if $j \in [1\,, \ell]$ then $(\underline{z}')_j = \underline{\tilde{z}}_j \geqslant z_j \geqslant x_j$, and if $j \in [1\,, t]$ then $g_j((\underline{z}')_{\ell+j}) = \underline{\tilde{z}}_{h(j)} \geqslant z_{h(j)} \geqslant g_j(y_j)$, so $(\underline{z}')_{\ell+j} \geqslant y_j$, as g_j is ascending. This finishes the proof of the assertion.

Since $\Gamma \subseteq \Lambda'$, we know that there exists a uniform local deterministic (respectively, randomized Las Vegas) algorithm $\mathcal{A}$ for Π and $\mathcal{F}$ such that the (respectively, expected) running time of $\mathcal{A}$ over any configuration $(G,\mathbf{x}) \in \mathcal{F}$ is $O(f'^* \cdot s_{f'}(f'^*)) = O(f'^* \cdot s_f(f'^*))$, where $f'^* = f(\mathtt{q}_1^*, \ldots, \mathtt{q}_\ell^*, \mathtt{p}_{r'+1}^*, \ldots, \mathtt{p}_r^*\})$. The fact that f' is non-decreasing implies that

$$f'^* \leqslant f'(\mathtt{q}_1^*, \ldots, \mathtt{q}_\ell^*, g_1^{-1}(\mathtt{q}_{h(1)}^*), \ldots, g_t^{-1}(\mathtt{q}_{h(t)}^*)) = f^*.$$

As s_f is non-decreasing, the (respectively, expected) running time of $\mathcal{A}$ is bounded by $O(f^* \cdot s_f(f^*))$, as desired. □

Applying Theorem 3 to the work of Barenboim and Elkin [6] (see Theorem 6.3 therein) with $\Gamma = \{a, n\}$ and $\Lambda = \{n\}$ yields the following result, since $a \leqslant n$.

Corollary 4 *The following uniform deterministic algorithms solving* `MIS` *exist :*

- *For the family of graphs with arboricity $a = o(\sqrt{\log n})$, running in time $o(\log n)$,*
- *For any constant $\delta \in (0\,, 1/2)$, for the family of graphs with arboricity $a = O(\log^{1/2-\delta} n)$, running in time $O(\log n / \log\log n)$.*

4.6 Running as Fast as the Fastest Algorithm

To illustrate the topic of the next theorem, consider the non-uniform algorithms for MIS for general graphs, namely, the algorithms of Barenboim and Elkin [4] and that of Kuhn [22], which run in time $O(\Delta + \log^* n)$ and use the knowledge of n and Δ, and the algorithm of Panconesi and Srinivasan [34], which runs in time $2^{O(\sqrt{\log n})}$ and requires the knowledge of n. Furthermore, consider the MIS algorithms of Barenboim and Elkin in [5,6], which are very efficient for graphs with a small arboricity a. If n, Δ and a are contained in the inputs of all nodes, then one can compare the running times of these algorithms and use the fastest one. That is, there exists a non-uniform algorithm $\mathcal{A}^{\{n,\Delta,a\}}$ that runs in time $T(n, \Delta, a) = \min\{g(n), h(\Delta, n), f(a, n)\}$, where $g(n) = 2^{O(\sqrt{\log n})}$, $h(\Delta, n) = O(\Delta + \log^* n)$, and $f(a, n)$ is defined as follows: $f(a, n) = o(\log n)$ for graphs of arboricity $a = o(\sqrt{\log n})$, $f(a, n) = O(\log n / \log\log n)$ for arboricity $a = O(\log^{1/2-\delta} n)$, for some constant $\delta \in (0\,, 1/2)$; and otherwise: $f(a, n) = O(a + a^\epsilon \log n)$, for arbitrary small constant $\epsilon > 0$.

Unfortunately, the theorems established so far do not allow us to transform $\mathcal{A}^{\{n,\Delta,a\}}$ into a uniform algorithm running in time $T(n, \Delta, a)$, the reason being that the function $T(n, \Delta, a)$ bounding the running time does not have a sequence number. On the other hand, as mentioned in Corollary 2, Theorem 1 does allow us to transform each of the algorithms in [4, 22,34] into a uniform MIS algorithm, with time complexity $O(\Delta + \log^* n)$ and $2^{O(\sqrt{\log n})}$, respectively. Moreover, Corollaries 3 and 4 show that Theorems 1 and 3 allow us to transform the algorithms in [5,6] to uniform algorithms that (over the appropriate graph families), run as fast as the corresponding non-uniform algorithms . Nevertheless, unless n, Δ and a are provided as inputs to the nodes, it is not clear how to obtain from these transformed algorithms a uniform algorithm running in time $T(n, \Delta, a)$. The following theorem solves this problem.

Theorem 4 *Consider a problem Π and a family of instances $\mathcal{F}$. Let k be a positive integer and let $\Lambda_1, \ldots, \Lambda_k$ be k sets of non-decreasing parameters. Let $\mathcal{P}$ be a $(\Lambda_1 \cup \cdots \cup \Lambda_k)$-monotone pruning algorithm for Π and $\mathcal{F}$. For $i \in \{1, 2, \cdots, k\}$, consider a uniform algorithm $\mathcal{U}_i$ whose running time is bounded with respect to Λ_i by a function f_i. Then there is a uniform algorithm with running time $O(f_{\min})$, where $f_{\min} = \min\{f_1(\Lambda_1^*), \ldots, f_k(\Lambda_k^*)\}$.*

Proof Clearly, it is sufficient to prove the theorem for the case $k = 2$. The basic idea behind the proof of theorem above is to run in iterations, such that each iteration i consists of running the quadruple $(\mathcal{U}_1; \mathcal{P}; \mathcal{U}_2; \mathcal{P})$, where $\mathcal{U}_1$ and $\mathcal{U}_2$ are executed for precisely 2^i rounds each. Hence, a correct solution will be produced in Iteration $s = \lceil \log f_{\min} \rceil$ or before. Since each iteration i takes at most $O(2^i)$ rounds (recall that the running time of $\mathcal{P}$ is constant), the running time is $O(f_{\min})$.

Formally, we define a sequence of uniform algorithms $(\mathcal{A}_i)_{i \in \mathbf{N}}$ as follows. For $i \in \mathbf{N}$, set $\mathcal{A}_{2i+1} = \hat{\mathcal{U}}_1$ and $\mathcal{A}_{2i+2} = \hat{\mathcal{U}}_2$, where $\hat{\mathcal{U}}_j$ is $\mathcal{U}_j$ restricted to 2^i rounds for $j \in \{1, 2\}$. Let π be the uniform alternating algorithm with respect to $(\mathcal{A}_i)_{i \in \mathbf{N}}$ and $\mathcal{P}$, that is $\pi = \mathcal{B}_1; \mathcal{B}_2; \mathcal{B}_3; \cdots$ where $\mathcal{B}_{2i+j} = \hat{\mathcal{U}}_j; \mathcal{P}$ for every $i \in \mathbf{N}$ and every $j \in \{1, 2\}$. Letting T_0 be the running time of $\mathcal{P}$, the running time of $\mathcal{B}_i$ is at most $2^i + T_0$, for every $i \in \mathbf{N}$.

Consider an instance $(G, \mathbf{x}) \in \mathcal{F}$. For each $(\mathtt{p}, \mathtt{q}) \in \Lambda_1 \times \Lambda_2$, let $\mathtt{p}^* = \mathtt{p}(G, \mathbf{x})$ and $\mathtt{q}^* = \mathtt{q}(G, \mathbf{x})$. Algorithm $\mathcal{B}_i$ operates on the configuration $(G_i, \mathbf{x}_i)$. Let $\mathtt{p} \in \Lambda_1 \cup \Lambda_2$. Because $\mathcal{P}$ is monotone with respect to $\Lambda_1 \cup \Lambda_2$, it follows by induction on i that $\mathtt{p}^* \geqslant \mathtt{p}(G_i, \mathbf{x}_i)$. Hence, the running time of $\mathcal{U}_j$ over $(G_i, \mathbf{x}_i)$ is bounded from above by $f_j(\Lambda_j^*)$ for every $i \in \mathbf{N}$ and each $j \in \{1, 2\}$. Thus, $V(G_{2s+2}) = \emptyset$ for the smallest s such that $2^s \geqslant f_{\min}$. In other words, $\pi = \mathcal{B}_1; \mathcal{B}_2; \cdots; \mathcal{B}_{2s+1}$. Consequently, by Observation 3.4, Algorithm π correctly solves Π on $\mathcal{F}$ and, since $\mathcal{B}_i$ runs in at most $2^{\lceil i/2 \rceil} + T_0$ rounds, the running time of π is $O(2^s) = O(f_{\min})$, as asserted. □

Now, we can combine Theorem 4 with Corollaries 3 and 4, and establish a uniform algorithm for MIS that runs in time $f(a, n)$. Combining this algorithm with Corollary 2, and applying once more Theorem 4 yields Corollary 1(i).

5 Uniform Coloring Algorithms

In general, we could not find a way to directly apply our transformers (e.g., the one given by Theorem 3) for the coloring problem. The main reason is that we could not find an efficient pruning algorithm for the coloring problem. Indeed, consider for example the $O(\Delta)$-coloring problem. The checking property of a pruning

algorithm requires that, in particular, the nodes can locally decide whether they belong to a legal configuration. While locally checking that neighboring nodes have distinct colors is easy, knowing whether a color is in the required range, namely, $[1\,,O(\Delta)]$, seems difficult as the nodes do not know Δ. Moreover, the gluing property seems difficult to tackle also: after pruning a node with color c, none of its unpruned neighbors can be colored in color c. In other words, a correct solution on the non-pruned subgraph may not glue well with the pruned subgraph.

Nevertheless, we show in this section that several relatively general transformers can be used to obtain uniform coloring algorithms from non-uniform one. We focus on standard coloring problems in which the required number of colors is given as a function of Δ.

5.1 Uniform $(\Delta+1)$-coloring Algorithms

A standard trick (cf., [28,30]) allows us to transform an efficient (with respect to n and Δ) MIS algorithm for general graphs into one for $(\Delta+1)$-coloring (and, actually, to the more general *maximal coloring* problem defined by Luby [30]). The general idea is based on the observation that $(\Delta+1)$-colorings of G and maximal independent sets of $G' = G \times K_{\Delta+1}$ are in one-to-one correspondence. More precisely, and avoiding the use of Δ, the graph G' is constructed from G as follows. For each node $u \in V(G)$, take a clique C_u of size $\deg_G(u)+1$ with nodes $u_1, \ldots, u_{\deg_G(u)+1}$. Now, for each $(u,v) \in E(G)$ and each $i \in [1\,, 1+\min\{\deg_G(u), \deg_G(v)\}]$, let $(u_i, v_i) \in E(G')$. The graph G' can be constructed by a local algorithm without using any global parameter. It remains to observe the existence of a natural one-to-one correspondence between the maximal independent sets of G' and the $(\deg_G+1)$-colorings of G, that is, the colorings of G such that each node u is assigned a color in $[1\,, \deg_G(u)+1]$.

To see this, first consider a $(\deg_G+1)$-coloring c of G. Set

$$X = \{u_i \in V(G') \,:\, c(u) = i\}\,.$$

Then, no two nodes in X are adjacent in G'. Moreover, a node that does not belong to X has a neighbor in X since X contains a vertex from each clique C_u for $u \in V(G)$. Therefore, X is a MIS of G'.

Conversely, let X be a MIS of G'. We assert that X contains a node from every clique C_u for $u \in V(G)$. Indeed, suppose on the contrary that $X \cap V(C_u) = \emptyset$ for a node $u \in V(G)$. By the definition of a MIS, every vertex $u_i \in V(C_u)$ has a neighbor $v(u_i)$ that belongs to X. Since a clique can contain at most one node in X and $v(u_i) \neq v(u_j)$ whenever $i \neq j$, we deduce that at least $|C_u|$ cliques C_v with $v \neq u$ contain a node that has a neighbor in C_u. This contradicts the definition of G', since $|C_u| = \deg_G(u)+1$. Thus, setting $c(u)$ to be the index $i \in \{1, \ldots, \deg_G(u)+1\}$ such that $u_i \in X$ yields a $(\deg_G+1)$-coloring of G.

Therefore, we obtain Corollary 1(ii) as a direct consequence of Corollary 1(i).

5.2 Uniform Coloring with More than $\Delta+1$ Colors

We now aim to provide a transformer taking as input an efficient non-uniform coloring algorithm that uses $g(\Delta)$ colors (where $g(\Delta) > \Delta$) and produces an efficient uniform coloring algorithm that uses $O(g(\Delta))$ colors. We begin with the following definitions.

An *instance* for the coloring problem is a pair $(G, \mathbf{x})$ where G is a graph and $\mathbf{x}(v)$ contains a color $c(v)$ such that the collection $\{c(v) \,:\, v \in V(G)\}$ forms a coloring of G. (The color $c(v)$ can be the identity $\mathrm{Id}(v)$.) For a given family $\mathcal{G}$ of graphs, we define $\mathcal{F}(\mathcal{G})$ to be the collection of instances $(G, \mathbf{x})$ for the coloring problem, where $G \in \mathcal{G}$.

Many coloring algorithms consider the identities as colors, and relax the assumption that the identities are unique by replacing it with the weaker requirement that the set of initial colors forms a coloring. Given an instance $(G, \mathbf{x})$, let $m = m(G, \mathbf{x})$ be the maximal identity. Note that m is a graph-parameter.

Recall the $\lambda(\tilde{\Delta}+1)$-coloring algorithms designed by Barenboim and Elkin [4] and Kuhn [22] (which generalize the $O(\tilde{\Delta}^2)$-coloring algorithm of Linial [28]). We would like to point out that, in fact, everything works similarly in these algorithms if one replaces n with m. That is, these $\lambda(\tilde{\Delta}+1)$-coloring algorithms can be viewed as requiring m and Δ and running in time $O(\tilde{\Delta}/\lambda + \log^* \tilde{m})$. The same is true for the edge-coloring algorithms of Barenboim and Elkin [7].

The following theorem implies that these algorithms can be transformed into uniform ones. In the theorem, we consider two sets Γ and Λ of non-decreasing graph-parameters such that

(1) Γ is weakly-dominated by Λ; and
(2) $\Gamma \subseteq \{\Delta, m\}$.

Two such sets of parameters are said to be *related*. The notion of moderately-fast function (defined in Section 2) will be used to govern the number of colors used by the coloring algorithms.

Theorem 5 *Let Γ and Λ be two related sets of non-decreasing graph-parameters and let $\mathcal{A}^{\Gamma}$ be a $g(\tilde{\Delta})$-coloring algorithm with running time bounded with respect to Λ by some function f. If*

1. *there exists a sequence-number function s_f for f;*
2. *g is moderately-fast;*
3. *the dependence of f on m is bounded by a polylog; and*
4. *the dependence of f on Δ is moderately-slow;*

then there exists a uniform $O(g(\Delta))$-coloring algorithm running in time $O(f(\Lambda^*) \cdot s_f(f(\Lambda^*)))$.

Proof Our first goal is to obtain a coloring algorithm that does not require m (and thus requires only Δ). For this purpose we first define the following problem.
The *strong list-coloring* (SLC) problem: a configuration for the SLC problem is a pair $(G, \mathbf{x}) \in \mathcal{F}(\mathcal{G})$ such that

(1) there exists an integer $\hat{\Delta}$ in $\cap_{v \in V(G)} \mathbf{x}(v)$ such that $\hat{\Delta} \geqslant \Delta$; and
(2) the input $\mathbf{x}(v)$ of every vertex $v \in V(G)$ contains a list $L(v)$ of colors contained in $[1\,, g(\hat{\Delta})] \times [1\,, \hat{\Delta}+1]$ such that

$$\forall k \in [1\,, g(\hat{\Delta})], \quad |\{j \,:\, (k,j) \in L(v)\}| \geqslant \deg_G(v) + 1.$$

Given a configuration $(G, \mathbf{x}) \in \mathcal{F}(\mathcal{G})$, an output vector $\mathbf{y}$ is a *solution to SLC* if it forms a coloring and if $\mathbf{y}(v) \in L(v)$ for every node $v \in V(G)$. Condition (1) above implies that a local algorithm for SLC can use an upper bound on Δ, which is the same for all nodes. Informally, Condition (2) above implies that the list $L(v)$ of colors available for each node v contains $\deg_G(v)+1$ copies of each color in the range $[1\,, g(\hat{\Delta})]$.

We now design a pruning algorithm $\mathcal{P}$ for SLC. Consider a triplet $(G, \mathbf{x}, \hat{\mathbf{y}})$, where $(G, \mathbf{x})$ is a configuration for SLC and $\hat{\mathbf{y}}$ is some tentative assignment of colors. The set W of nodes to be pruned is composed of all nodes u satisfying $\hat{\mathbf{y}}(u) \in L(u)$ and $\hat{\mathbf{y}}(u) \neq \hat{\mathbf{y}}(v)$ for all $v \in N_G(u)$. For each node $u \in V \setminus W$, set

$$L'(u) = L(u) \setminus \{\hat{\mathbf{y}}(v) \,:\, v \in N_G(u) \cap W\}\,.$$

In other words, $L'(u)$ contains all the colors in $L(u)$ that are not assigned to a neighbor of u belonging to W. Algorithm $\mathcal{P}$ returns the configuration $(G', \mathbf{x}')$, where G' is the subgraph of G obtained by removing the nodes in W and

$$\mathbf{x}'(u) = (\mathbf{x}(u) \setminus L(u)) \cup L'(u), \quad \text{for } u \in V \setminus W.$$

Observe that if we start with a configuration $(G, \mathbf{x})$ for SLC, then the output $(G', \mathbf{x}')$ of the pruning algorithm $\mathcal{P}$ is also a configuration for SLC. Indeed, for every node v and every integer k, at most $\deg_W(v)$ pairs (k, j) are removed from the list $L(v)$ of v, where $\deg_W(v)$ is the number of neighbors of v that belong to W. On the other hand, the degree of v in G' is reduced by $\deg_W(v)$. Note also that the input vector of all nodes still contain $\hat{\Delta}$, which is an upper bound for the maximum degree of G'.

Starting with $\mathcal{A}^{\Gamma}$, it is straightforward to design a local algorithm $\mathcal{B}^{\Gamma'}$ for *SLC* that depends on $\Gamma' = \Gamma \setminus \{\Delta\}$. Specifically, $B^{\Gamma'}$ executes A^{Γ} using the good guess $\tilde{\Delta} = \hat{\Delta}$ for the parameter Δ. Furthermore, if A^{Γ} outputs at v a color c, then $B^{\Gamma'}$ outputs the color (c, j) where $j = \min\{s \,:\, (c,s) \in L(v)\}$.

Given an instance for SLC, we view $\hat{\Delta}$ as a non-decreasing parameter, and convert Λ to a new set of non-decreasing parameters Λ' by replacing Δ with Δ'. Formally, if $\Delta \in \Lambda$ then set $\Lambda' = (\Lambda \setminus \Delta) \cup \hat{\Delta}$, and otherwise, set $\Lambda' = \Lambda$. Since Γ and Λ contain only non-decreasing graph-parameters—and since $\hat{\Delta}$ is contained in all the inputs—we deduce that the pruning algorithm $\mathcal{P}$ is $(\Gamma' \cup \Lambda')$-monotone.

Now, we apply Theorem 3 to Algorithm $\mathcal{B}^{\Gamma'}$, the sets Γ' and Λ' of non-decreasing parameters and the aforementioned pruning algorithm $\mathcal{P}$ for SLC. We obtain a uniform algorithm $\mathcal{B}$ for SLC and $\mathcal{F}(\mathcal{G})$, whose running time is $O(f(\Lambda'^*) \cdot s_f(f(\Lambda'^*)))$.

We are ready to specify the desired uniform $O(g(\Delta))$-coloring algorithm. We define inductively a sequence $(D_i)_{i \in \mathbf{N}}$ by setting $D_1 = 1$ and

$$D_{i+1} = \min\{\ell \,:\, g(\ell) \geqslant 2g(D_i)\}$$

for $i \geqslant 1$. As g is moderately-increasing, there is a constant α such that for each integer $i \geqslant 1$,

1. $D_{i+1} \geqslant \alpha D_i$ and
2. $g(D_{i+1}) \leqslant \alpha^{\log \alpha} g(D_i)$.

Given an initial configuration $(G, \mathbf{x})$, we partition it according to the node degrees. For $i \in \mathbf{N}$, let G_i be the subgraph of G induced by the set of nodes v of G with $\deg_G(v) \in [D_i\,, D_{i+1}-1]$. Let $\mathbf{x}_i$ be the input $\mathbf{x}$ restricted to the nodes in G_i. The configuration $(G_i, \mathbf{x}_i)$, which belongs to $\mathcal{F}(\mathcal{G})$, is referred to as *layer i*. Note that nodes can figure out locally which layer they belong to. Observe also that $D_{i+1}-1$ is an upper bound on node degrees in layer i.

The algorithm proceeds in two phases. In the first phase, each node in layer i is assigned the list of colors $L''_i = [1\,, g(D_{i+1})] \times [1\,, D_{i+1}+1]$, and the degree estimation $\hat{\Delta}_i = D_{i+1}$. Each layer is now an instance of SLC and we execute Algorithm $\mathcal{B}$ in parallel on all layers. If Algorithm $\mathcal{B}$ assigns a color (c, j) to a node v in layer i then we change this color to $(g(D_{i+1}) + c, j)$. Hence, for each i, layer i is colored with colors taken from $L'_i = [g(D_{i+1})+1\,, 2g(D_{i+1})] \times [1\,, D_{i+1}+1]$.

Note that nodes in different layers have disjoint color lists, and hence we obtain a coloring of the whole graph G. The number of colors in L'_i is at most $2D_{i+1}g(D_{i+1})$.

Let $i_{\max}$ is the maximal integer i such that layer i is non-empty. The total number of colors used in the first phase is at most $2D_{i_{\max}+1}g(D_{i_{\max}+1})$, which is $O(\Delta g(\Delta))$ by Properties 1 and 2 above.

Furthermore, the running time of the first phase of the algorithm is dominated by the running time of the algorithm on layer $i_{\max}$. That is, the running time is at most $O(f(\Lambda'^*) \cdot s_f(f(\Lambda'^*)))$, where Λ'^* is the collection of correct parameters in Λ' for layer $i_{\max}$. Since $D_{i_{\max}+1} = O(\Delta)$ and the dependence of f on Δ is moderately-slow, we infer that $f(\Lambda'^*) = O(f(\Lambda^*))$. As s_f is moderately-slow too (by the definition), we deduce that the running time is $O(f(\Lambda^*) \cdot s_f(f(\Lambda^*)))$.

The second phase consists of running a second algorithm to change the set of possible colors of nodes in layer i from L'_i to $L_i = [g(D_{i+1}) + 1\,, 2g(D_{i+1})]$. Specifically, on layer i, we execute $\mathcal{A}^\Gamma$ using the guess $\tilde{\Delta} = D_{i+1}$ for the parameter Δ and the guess $\tilde{m} = 2D_{i+1}g(D_{i+1})$ for the parameter m (recall that $\Gamma \subseteq \{\Delta, m\}$). This procedure colors each layer with colors taken from the range $[1\,, g(D_{i+1})]$. Let v be in layer i and let $c(v)$ be the color assigned to v by $\mathcal{A}^\Gamma$. The final color of v given by our desired algorithm $\mathcal{A}$ is $g(D_{i+1}) + c(v)$. Thus, the colors assigned to the nodes in layer i belong to $[g(D_{i+1}) + 1\,, 2g(D_{i+1})]$. Therefore, nodes in different layers are assigned distinct colors. The algorithm is executed on each layer independently, all in parallel. Hence, we obtain a coloring of G. Moreover, since g is moderately-increasing, the total number of colors used is $O(g(\Delta))$.

Recall that $D_{i+1} = O(\Delta)$ and $g(D_{i+1}) = O(g(\Delta))$ for all i such that G_i is not empty. Hence, we deduce that the running time of the second phase of the algorithm is bounded from above by the running time of $\mathcal{A}^\Gamma$ on $(G, \mathbf{x})$ using the guesses $\tilde{\Delta} = O(\Delta)$ and $\tilde{m} = O(\Delta g(\Delta))$. Moreover, the fact that $g(x)$ is bounded by a polynomial in x implies that $\tilde{m}$ is at most polynomial in Δ, and hence in m.

Now, as the dependence of f on Δ is moderately-slow and the dependence of f on m is polylogarithmic, the running time of the second phase of $\mathcal{A}$ is $O(f(\Lambda))$. Combining this with the running time of the first phase concludes the proof. □

By Observation 4.1, the constant function $s_f = 1$ is a sequence-number function for every additive function f. Hence, Corollary 1(iii) directly follows from Theorem 5. Regarding edge-coloring, observe that Barenboim and Elkin [7] obtain their algorithm for general graphs by running a vertex-coloring algorithm $\mathcal{A}$ on the line-graph of the given graph. This algorithm $\mathcal{A}$ uses m and Δ in and the number of colors and time complexity of the resulting edge-coloring algorithm are that of $\mathcal{A}$. Using Theorem 5, one can transform the algorithm $\mathcal{A}$ designed for the family of line graphs into a uniform one, having asymptotically the same number of colors and running time. Hence, Theorem 1(v) follows.

Let $f\colon \mathbf{N}^2 \to \mathbf{R}$ be given by $f(x_1, x_2) = f_1(x_1) \cdot f_2(x_2)$, where f_1 and f_2 are ascending functions. By Observation 4.1, the function $s_f(i) = \lceil \log i \rceil + 1$ is a sequence-number function for f. Therefore, Corollary 1(iv) now follows by applying Theorem 5 to the coloring algorithms of Barenboim and Elkin [5].

6 Conclusion and Further Research

6.1 Pruning Algorithms

This paper focuses on removing assumptions concerning global knowledge in the context of local algorithms. We provide transformers taking a non-uniform local algorithm as a black box and producing a uniform algorithm running in asymptotically the same number of rounds. This is established via the notion of pruning algorithms. We believe that this novel notion is of independent interest and can be used for other purposes too, e.g., in the context of fault tolerance or dynamic settings.

We remind the reader that we restricted the running time of a pruning algorithm to be constant. This is because in all our applications we use constant time pruning algorithms. In fact, our transformers extend to the case where the given *uniform* pruning algorithm $\mathcal{P}$

- has running time bounded with respect to a set $\mathcal{S}$ of non-decreasing parameters by a (non-decreasing) function h; and
- is $\mathcal{S}$-monotone.

However, the transformer may incur an additive overhead in the running time of the obtained uniform algorithms, as these repeatedly use $\mathcal{P}$. Specifically, the overhead will be $h(\mathcal{S}^*)$ times the number of iterations used by the transformer, which is typically logarithmic in the running time of the non-uniform algorithm. It would be interesting to have an example of a problem that admits a fast non-trivial uniform pruning algorithm but does not admit a constant time one.

6.2 Bounded Message Size

This paper focuses on the $\mathcal{LOCAL}$ model, which does not restrict the number of bits used in messages. Ideally, messages should be short, i.e., using $O(\log n)$ bits. We found it difficult to obtain a general transformer that takes an arbitrary non-uniform algorithm using short

messages and produces a uniform one having asymptotically the same running time and message size. The reason is that techniques similar to those used in this paper, require guesses that fit for both the function bounding the running time and the function bounding the message size. Nevertheless, maintaining the same message size may still be possible given particular non-uniform algorithms that use messages whose content does not depend on the guessed upper bounds, such as algorithms that encode in the messages only identifiers, colors, or degrees.

6.3 Coloring

Recall that one of the difficulties in obtaining a pruning algorithm for coloring problems lies in the fact that the gluing property may not hold, that is, a pruned node v with color c may have a non-pruned neighbor u which is also colored c in some correct coloring of the non-pruned subgraph. In the context of running in iterations, in which one invokes a pruning algorithm and subsequently, an algorithm $\mathcal{A}$ on the non-pruned subgraph (similarly to Theorem 3), the aforementioned undesired phenomenon could be prevented if the algorithm $\mathcal{A}$ would avoid coloring node u with color c. With this respect, we believe that it would be interesting to investigate connections between g-coloring problems and *strong* g-coloring problems, in which each node v is given as an input a list of (forbidden) colors $F(v)$. In a correct solution, each node v must color itself with a color not in $L(v)$ such that the final configuration is a coloring using at most g colors.

Finally, recall that our transformer for coloring applies to deterministic algorithms only. It would be interesting to design a general transformer that takes non-uniform randomized coloring algorithms (e.g., the ones by Schneider and Wattenhofer [36]) and transforms them into uniform ones with asymptotically the same running time.

Acknowledgements The authors thank Boaz Patt-Shamir and the anonymous referees for their careful reading and thoughtful suggestions. Their comments helped to considerably improve the presentation of the paper.

References

1. Alon, N., Babai, L., Itai, A.: A fast and simple randomized parallel algorithm for the maximal independent set problem. J. Algorithms **7**, 567–583 (1986)
2. Awerbuch, B.: Complexity of network synchronization. J. ACM **32**, 804–823 (1985)
3. Awerbuch, B., Luby, M., Goldberg, A.V., Plotkin, S.A.: Network decomposition and locality in distributed computation. In: Proc. 30th IEEE Symp. Found. Comput. Sci. (FOCS), pp. 364–369 (1989)
4. Barenboim, L., Elkin, M.: Distributed $(\Delta + 1)$-coloring in linear (in Δ) time. In: Proc. 41st ACM Symp. Theor. Comput. (STOC), pp. 111–120 (2009)
5. Barenboim, L., Elkin, M.: Deterministic distributed vertex coloring in polylogarithmic time. In: Proc. 29th ACM Symp. Principles Distrib. Comput. (PODC), pp. 410–419 (2010)
6. Barenboim, L., Elkin, M.: Sublogarithmic distributed mis algorithm for sparse graphs using nash-williams decomposition. Distrib. Comput. **22**(5-6), 363–379 (2010)
7. Barenboim, L., Elkin, M.: Distributed deterministic edge coloring using bounded neighborhood independence. In: Proc. 30th ACM Symp. Principles Distrib. Comput. (PODC) (2011)
8. Bentley, J.L., Yao, A.C.C.: An almost optimal algorithm for unbounded searching. Information Processing Lett. **5**(3), 82–87 (1976)
9. Cohen, R., Fraigniaud, P., Ilcinkas, D., Korman, A., Peleg, D.: Label-guided graph exploration by a finite automaton. ACM Trans. Algorithms **4**, 42:1–42:18 (2008)
10. Cole, R., Vishkin, U.: Deterministic coin tossing and accelerating cascades: micro and macro techniques for designing parallel algorithms. In: Proc. 18th ACM Symp. Theor. Comput. (STOC), pp. 206–219 (1986)
11. Derbel, B., Gavoille, C., Peleg, D., Viennot, L.: On the locality of distributed sparse spanner construction. In: Proc. 27th ACM Symp. Principles Distrib. Comput. (PODC), pp. 273–282 (2008)
12. Dereniowski, D., Pelc, A.: Drawing maps with advice. In: Proc. 24th Int. Symp. Distrib. Comput. (DISC), pp. 328–342. Springer-Verlag (2010)
13. Fraigniaud, P., Gavoille, C., Ilcinkas, D., Pelc, A.: Distributed computing with advice: information sensitivity of graph coloring. Distrib. Comput. **21**, 395–403 (2009)
14. Fraigniaud, P., Ilcinkas, D., Pelc, A.: Communication algorithms with advice. J. Comput. Syst. Sci. **76**, 222–232 (2010)
15. Fraigniaud, P., Korman, A., Lebhar, E.: Local mst computation with short advice. In: Proc. 19th ACM Symp. Parallelism Algo. Archit. (SPAA), pp. 154–160 (2007)
16. Fraigniaud, P., Korman, A., Peleg, D.: Local distributed decision. Submitted for Publication
17. Goldberg, A.V., Plotkin, S.A.: Efficient parallel algorithms for $(\Delta+1)$-coloring and maximal independent set problem. In: Proc. 19th ACM Symp. Theor. Comput. (STOC), pp. 315–324 (1987)
18. Goldberg, A.V., Plotkin, S.A., Shannon, G.E.: Parallel symmetry-breaking in sparse graphs. SIAM J. Discrete Math. **1**(4), 434–446 (1988)
19. Hańćkowiak, M., Karoński, M., Panconesi, A.: On the distributed complexity of computing maximal matchings. SIAM J. Discrete Math. **15**(1), 41–57 (electronic) (2001/02)
20. Korman, A., Kutten, S.: Distributed verification of minimum spanning trees. Distrib. Comput. **20**, 253–266 (2007)
21. Korman, A., Kutten, S., Peleg, D.: Proof labeling schemes. Distrib. Comput. **22**, 215–233 (2010)
22. Kuhn, F.: Weak graph colorings: distributed algorithms and applications. In: Proc. 21st ACM Symp. Parallelism Algo. Archit. (SPAA), pp. 138–144 (2009)
23. Kuhn, F., Moscibroda, T., Wattenhofer, R.: What cannot be computed locally! In: Proc. 23rd ACM Symp. Principles Distrib. Comput. (PODC), pp. 300–309 (2004)

24. Kuhn, F., Wattenhofer, R.: On the complexity of distributed graph coloring. In: Proc. 25th ACM Symp. Principles Distrib. Comput. (PODC), pp. 7–15 (2006)
25. Kutten, S., Peleg, D.: Tight fault locality. SIAM J. Comput. **30**(1), 247–268 (electronic) (2000)
26. Lenzen, C., Oswald, Y., Wattenhofer, R.: What can be approximated locally?: case study: dominating sets in planar graphs. In: Proc. 20th ACM Symp. Parallelism Algo. Archit. (SPAA), pp. 46–54 (2008)
27. Linial, N.: Distributive graph algorithms global solutions from local data. In: Proc. 28th IEEE Symp. Found. Comput. Sci. (FOCS), pp. 331–335 (1987)
28. Linial, N.: Locality in distributed graph algorithms. SIAM J. Comput. **21**, 193 (1992)
29. Lotker, Z., Patt-Shamir, B., Rosén, A.: Distributed approximate matching. SIAM J. Comput. **39**(2), 445–460 (2009)
30. Luby, M.: A simple parallel algorithm for the maximal independent set problem. SIAM J. Comput. **15**, 1036–1053 (1986)
31. Nakano, K., Olariu, S.: Uniform leader election protocols for radio networks. IEEE Trans. Parallel Distrib. Syst. **13**(5), 516–526 (2002)
32. Naor, M., Stockmeyer, L.: What can be computed locally? SIAM J. Comput. **24**(6), 1259–1277 (1995)
33. Panconesi, A., Rizzi, R.: Some simple distributed algorithms for sparse networks. Distrib. Comput. **14**, 97–100 (2001)
34. Panconesi, A., Srinivasan, A.: On the complexity of distributed network decomposition. J. Algorithms **20**(2), 356–374 (1996)
35. Peleg, D.: Distributed computing. A locality-sensitive approach. SIAM Monographs on Discrete Mathematics and Applications, 5. Philadelphia, PA: SIAM, Society for Industrial and Applied Mathematics. xvi, 343 p. (2000)
36. Schneider, J., Wattenhofer, R.: A new technique for distributed symmetry breaking. In: Proc. 29th ACM Symp. Principles Distrib. Comput. (PODC), pp. 257–266 (2010)
37. Schneider, J., Wattenhofer, R.: An optimal maximal independent set algorithm for bounded-independence graphs. Distrib. Comput. **22**(5-6), 1–13 (2010)
38. Szegedy, M., Vishwanathan, S.: Locality based graph coloring. In: Proc. 25th ACM Symp. Theor. Comput. (STOC), pp. 201–207 (1993)